



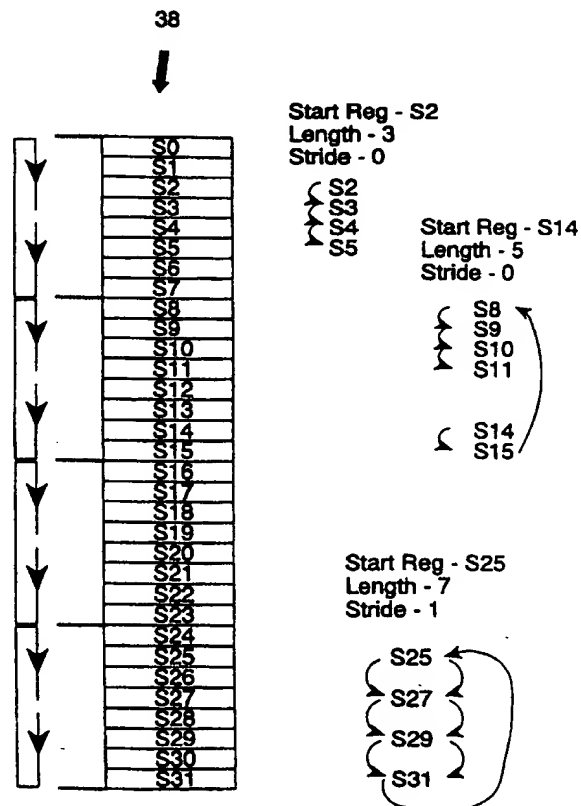
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 15/78, 9/30		(11) International Publication Number: WO 99/61997
A1		(43) International Publication Date: 2 December 1999 (02.12.99)
(21) International Application Number: PCT/GB99/00707		(81) Designated States: CN, IL, IN, JP, KR, RU, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).
(22) International Filing Date: 9 March 1999 (09.03.99)		
(30) Priority Data: 09/085,752 27 May 1998 (27.05.98) US		
(71) Applicant: ARM LIMITED [GB/GB]; 90 Fulbourn Road, Cherry Hinton, Cambridge CB1 4JN (GB).		
(72) Inventors: HINDS, Christopher, Neal; 6400 Back Bay Lane, Austin, TX 78739 (US). JAGGAR, David, Vivian; 5704 Galsworthy Court, Austin, TX 78739 (US). MATHENY, David, Terrence; 10706 Buckthorn Drive, Austin, TX 78759 (US). SEAL, David, James; 68 Keynes Road, Cambridge CB5 8PR (GB).		
(74) Agent: ROBINSON, Nigel, Alexander, Julian; D. Young & Co., 21 New Fetter Lane, London EC4A 1DA (GB).		Published With international search report.

(54) Title: RECIRCULATING REGISTER FILE

(57) Abstract

A floating point unit having a register bank containing a plurality of registers supports vector operations that execute a specified operation a plurality of times upon a sequence of data values from different registers. The register bank is divided into subsets and with the sequence of registers used in a vector operation wrapping within a subset. The subsets comprise disjoint, contiguous ranges of register numbers. The wrapping within ranges allows compact code and efficient to be provided for performing DSP operations, such as FIR filtering and matrix transformations.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon	KR	Republic of Korea	PL	Poland		
CN	China	KZ	Kazakstan	PT	Portugal		
CU	Cuba	LC	Saint Lucia	RO	Romania		
CZ	Czech Republic	LI	Liechtenstein	RU	Russian Federation		
DE	Germany	LK	Sri Lanka	SD	Sudan		
DK	Denmark	LR	Liberia	SE	Sweden		
EE	Estonia			SG	Singapore		

RECIRCULATING REGISTER FILE

This invention relates to the field of data processing. More particularly, this invention relates to data processing systems having a register bank and supporting vector operations.

5 It is known to provide data processing systems having a register bank and supporting vector operations. Examples of such systems are the Cray 1 and Digital Equipment Corporation MultiTitan processors.

The Cray 1 processor has separate vector and scalar register banks. If the opcode of the instruction being executed indicates a vector operation, then a sequence of
10 data values are returned from the vector register bank in dependence upon a length value stored in a length register and a mask stored in a mask register. The length specifies how many data values are in the sequence and mask specifies which data values are returned from among a plurality of data values associated with the vector register indicated in the instruction.

15 The MultiTitan processor has a single register bank the registers of which can act serve either as scalars or vectors. The instruction itself includes flags that indicate whether a register specified is a scalar or a vector and a length field indicating the number of data values within the sequence when a vector register is used.

Vector instructions themselves are desirable as they allow code density to be
20 increased since a single instruction can specify a plurality of data processing operations. Digital signal processing such as audio or graphics processing is particularly well suited to exploiting vector operations as there is often a requirement to perform the same operation upon a sequence of related data values, e.g. performing a filter operation by multiply a sequence of signal values by tap coefficients of a digital filter.

25 It is also desirable to perform data processing operations as quickly and efficiently as possible. One way of helping to increase speed and efficiency is to avoid having to reload or reposition data values that have already been stored within the register bank. A problem in achieving this is that instruction code that is able to reuse data values directly tends to be longer and more complex. If more instructions are

needed to specify the operation required then this tends to slow down the processing and negate the purpose of seeking to reuse the data values within the register bank.

As an alternative to the use of general purpose processors such as the Cray 1 and MultiTitan processors, special purpose digital signal processing circuits are often
5 provided with the specific role of supporting a small number of digital signal processing operations. Within these special purpose digital signal processing circuits, a common technique is to store the data values required within a large memory and then fetch the data values required for each manipulation as needed. The data values need not be reloaded or reposition within the large memory as the order and sequencing of their use
10 is controlled by manipulation of the addresses used to access the large memory. A problem with this approach is that the circuits have to be specifically designed to match the operation being performed and so lack the flexibility and ease of integration with other functions that is provided by the use of a more typical general purpose processor.

It is an object of the present invention to provide efficient and fast data
15 processing whilst maintaining the flexibility of a general purpose processor using a register bank and instruction decoder supporting vector operations.

Viewed from one aspect the present invention provides an apparatus for processing data, said apparatus comprising:

a register bank having a plurality of registers; and
20 an instruction decoder responsive to at least one data processing instruction specifying a vector operation that executes a data processing operation a plurality of times using data values from a sequence of registers within said register bank; wherein

said register bank includes at least one subset of registers, said sequence of
25 registers being within said subset; and
said instruction decoder controls said sequence of registers to wrap within said subset of registers.

Providing register wrapping within a subset (i.e. less than all) of the registers of the register bank allows compact code to be written that reuses data values within

the register bank without requiring reloading or moving of the data values. The instruction codes for each use can start at different points within the subset and so act upon the data values in a different order with the hardware taking care of the wrapping needed without having to provide extra instructions to split the vector sequence.

- 5 Furthermore, performing a vector operation upon a subset of registers that wrap upon themselves allows for the possibility to concurrently perform data transfers upon data values within registers not within the subset. The register wrapping can also be thought of as providing hardware support for a ring (circular) buffer type arrangement where, for example, data is loaded in and multiplied out of the buffer at point that
- 10 chase each other round and round the buffer.

Whilst it is possible to have only a single subset of wrapping registers, it is advantageous to provide systems in which said vector operation executes said data processing operation using a plurality of respective data values from a corresponding plurality of sequences of registers;

- 15 said register bank contains a plurality of subsets of registers, said plurality of sequences of registers being within respective subsets; and

said instruction decoder controls said plurality of sequences of registers to wrap within respective subsets of registers.

- Within digital signal processing operations it is often required to reuse the data
- 20 values from two sequences (e.g. a FIR operation with taps and signal values to be multiplied and accumulated at different offsets or a matrix operation) and so multiple subsets of wrapping registers are desirable.

- Whilst it is possible for the subsets to overlap, in practice the data values that require reuse in such situations are usually quite separate and so the subsets may be
- 25 made disjoint. This makes the hardware implementation advantageously less complex.

It will be appreciated that the subsets could be formed of registers from positions intermixed with registers not within the subsets. However, programming and implementation are made easier when the subsets are a range of consecutively

numbered registers.

The ranges could be spaced apart within the register bank, but in preferred embodiments the ranges are contiguous as this makes more efficient use of the register space available.

- 5 The ability of the invention to more effectively use vector operations is complemented in preferred embodiments which further comprise a memory and a transfer controller for controlling transfers of data values between said memory and registers within said register bank, said transfer controller being responsive to multiple transfer instructions to transfer a sequence of data values between said
- 10 memory and a sequence of registers within said register bank.

The ability to transfer data values to and from blocks of registers within the register bank is well matched to the ability of the invention to efficiently use vector operations as it allows a block of registers to be reused several times and then swapped out with a single instruction.

- 15 The use of a sequence of registers in the vector operation and the division of the register bank into predefined subsets of registers may be efficiently implemented in preferred embodiments in which each range is addressed via an incrementer that wraps between the end points of that range.

- The sequence of registers used in the vector operation could take many forms,
- 20 e.g. every alternate register within the subset, however the most commonly useful situation is the one in which the sequence is a sequence of consecutive registers.

- The above techniques can be used in any processor have a register bank and supporting vector operations. However, the ability to allow compact code and reuse data values within the registers has been found to be particularly useful and not interfere
- 25 with other considerations in embodiments in which the register bank and instruction decoder are within a floating point unit.

Viewed from another aspect the invention provides a method of processing data, said method comprising the steps of:

storing data values within a plurality of registers of a register bank; and

in response to at least one data processing instruction specifying a vector operation, executing a data processing operation a plurality of times using data values from a sequence of registers within said register bank; wherein

said register bank includes at least one subset of registers, said sequence of
5 registers being within said subset; and

during said executing, said sequence of registers wraps within said subset of registers.

The technique is particularly useful in efficiently providing FIR filter operations in which tap coefficient values and signal values are reused several times
10 with the relative offset between them changing for each vector operation.

Embodiments of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1 schematically illustrates a data processing system;

Figure 2 illustrates a floating point unit supporting both scalar and vector
15 registers;

Figure 3 is a flow diagram illustrating how, for single precision operation, it is determined whether a given register is a vector or scalar register;

Figure 4 is a flow diagram illustrating how, for double precision operation, it is determined whether a given register is a vector or a scalar;

20 Figure 5 illustrates the division of the register bank into subsets with wrapping within each subset during single precision operation;

Figure 6 illustrates the division of the register bank into subsets with wrapping within each subset during double precision operation;

Figures 7A to 7C illustrate a main processor view of a coprocessor instruction, a
25 single and double precision coprocessor view of the coprocessor instruction and a single precision coprocessor view of the coprocessor instruction respectively;

Figure 8 illustrates a main processor controlling a single and double precision coprocessor;

Figure 9 illustrates the main processor controlling a single precision

coprocessor;

Figure 10 illustrates the circuit within the single and double precision coprocessor that determines whether an accept signal should be returned to the main processor for a received coprocessor instruction;

- 5 Figure 11 illustrates the circuit within the single precision coprocessor that determines whether an accept signal should be returned to the main processor for a received coprocessor instruction;

Figure 12 illustrates undefined instruction exception handling within the main processor;

- 10 Figure 13 is a block diagram illustrating elements of a coprocessor in accordance with preferred embodiments of the present invention;

Figure 14 is a flow diagram illustrating operation of the register control and instruction issue logic in accordance with preferred embodiments of the present invention;

- 15 Figure 15 provides an example of the contents of the floating point register in accordance with preferred embodiments of the present invention;

Figure 16 illustrates the register bank within a Cray 1 processor; and

Figure 17 illustrates the register bank within a MultiTitan processor.

Figure 1 illustrates a data processing system 22 comprising a main processor 24, a floating point unit coprocessor 26, a cache memory 28, a main memory 30 and an input/output system 32. The main processor 24, the cache memory 28, the main memory 30 and the input/output system 32 are linked via a main bus 34. A coprocessor bus 36 links the main processor 24 to the floating point unit coprocessor 26.

In operation, the main processor 24 (also referred to as the ARM core) executes a stream of data processing instructions that control data processing operations of a general type including interactions with the cache memory 28, the main memory 30 and the input/output system 32. Embedded within the stream of data processing instructions are coprocessor instructions. The main processor 24 recognises these coprocessor instructions as being of a type that should be executed by an attached coprocessor. Accordingly, the main processor 24 issues these coprocessor instructions on the coprocessor bus 36 from where they are received by any attached coprocessors. In this case, the floating point unit coprocessor 26 will accept and execute any received coprocessor instructions that it detects are intended for it. This detection is via a coprocessor number field within the coprocessor instruction.

Figure 2 schematically illustrates the floating point unit coprocessor 26 in more detail. The floating point unit coprocessor 26 includes a register bank 38 that is formed of 32 32-bit registers (less shown in Figure 2). These registers can operate individually as single precision registers each storing a 32-bit data value or as pairs that together store a 64-bit data value. Within the floating point unit coprocessor 26 there is provided a pipelined multiply accumulate unit 40 and a load store control unit 42. In appropriate circumstances, the multiply accumulate unit 40 and the load store control unit 42 can operate concurrently with the multiply accumulate unit 40 performing arithmetic operations (that include multiply accumulate operations as well

as other operations) upon data values within the register bank 38 whilst the load store control unit 42 transfers data values not being used by the multiply accumulate unit 40 to and from the floating point unit coprocessor 26 via the main processor 24.

Within the floating point unit coprocessor 26, a coprocessor instruction that is
5 accepted is latched within an instruction register 44. The coprocessor instruction can in this simplified view be considered to be formed of an opcode portion followed by three register specifying fields R1, R2 and R3 (in fact these fields may be split and spread around differently within a full instruction). These register specifying fields R1, R2 and R3 respectively correspond to the registers within the register bank 38 that
10 serve as the destination, first source and second source for the data processing operation being performed. A vector control register 46 (which may be part of a larger register serving additional functions) stores a length value and a stride value for the vector operations that may be performed by the floating point unit coprocessor 26.

The vector control register 46 may be initialised and updated with length and stride
15 values in response to a vector control register load instruction. The vector length and stride values apply globally within the floating point unit coprocessor 26 thereby allowing these values to be dynamically altered on a global basis without having to resort to self-modifying code.

A register control and instruction issue unit 48, the load store control unit 42
20 and a vector control unit 50 can together be considered to perform a main part of the role of instruction decoder. The register control and instruction issue unit 48 is responsive to the opcode and the three register specifying fields R1, R2 and R3 and first outputs the initial register access (address) signals to the register bank 38 without performing any decode upon the opcode or needing to use the vector control unit 50.
25 Having direct access to the initial register values in this way assists in achieving a faster implementation. If a vector register is specified, then the vector control unit 50 serves to generate the necessary sequence of register access signals using 3-bit incrementers (adders) 52. The vector control unit 50 is responsive to the length value and the stride value stored within the vector control register 46 in performing its

addressing of the register bank 38. A register scoreboard 54 is provided to perform register locking such that the pipelined multiply accumulate unit 40 and concurrently operating load store control unit 42 do not give rise to any data consistency problems (the register scoreboard 54 may alternatively be considered to be part of the register control and instruction issue unit 48).

The opcode within the instruction register 44 specifies the nature of the data processing operation to be performed (e.g. whether the instruction is an add, subtract, multiply, divide, load, store... etc.). This is independent of the vector or scalar nature of the register being specified. This further simplifies the instruction decoding and set-up of the multiply accumulate unit 40. The first register specifying value R1 and the second register specifying value R2 together encode the vector/scalar nature of the operation specified by the opcode. The three common cases supported by the encoding are $S = S * S$ (e.g. basic random maths as generated by a C compiler from a block of C code), $V = V \text{ op } S$ (e.g. to scale the elements of a vector) and $V = V \text{ op } V$ (e.g. matrix operations such as FIR filters and graphics transformations) (note that in this context a "op" indicates a general operation and the syntax is of the form destination = second operand op first operand). It should also be understood that some instructions (e.g. a compare, a compare with zero or an absolute value) may have no destination registers (e.g. outputs are the condition flags) or fewer input operands (a compare with zero has only one input operand). In these cases there is more opcode bit space available to specify options such as vector/scalar nature and the full range of registers could be made available for each operand (e.g. compares may always be fully scalar whatever the register).

The register control and instruction issue unit 48 and the vector control unit 50 that together perform the main part of the role of instruction decoder are responsive to the first register specifying field R1 and the second register specifying field R2 to determine and then control the vector/scalar nature of the data processing operation specified. It will be noted that if the length value stored within the vector control register 46 indicates a length of one (corresponding to a stored value of zero), then

this can be used as an early indication of a purely scalar operation.

Figure 3 shows a flow diagram indicating the processing logic used to decode the vector/scalar nature from the register specifying values in single precision mode. At step 56 a test is made as to whether the vector length is globally set as one (length value equals zero). If the vector length is one, then all registers are treated as scalars in step 58. At step 60, a test is made as to whether the destination register R1 is within the range S0 to S7. If this is the case, then the operation is all scalar and is of the form $S = S \text{ op } S$, as is indicated in step 62. If step 60 returns a no, then the destination is determined to be a vector, as indicated at step 64. If the destination is a vector, then the encoding takes the second operand as also being a vector. Accordingly, the two possibilities remaining at this stage are $V = V \text{ op } S$ and $V = V \text{ op } V$. These two options are distinguished between by the test at step 66 that determines whether the first operand is one of S0 to S7. If this is the case, then the operation is $V = V \text{ op } S$, else the operation is $V = V \text{ op } V$. These states are recognised in steps 68 and 70 respectively.

It should be noticed that when the vector length is set to one, then all of the 32 registers of the register bank 38 are available to be used as scalars since the scalar nature of the operation will be recognised at step 58 without having to rely upon the test of step 60 that does limit the range of registers that may be used for the destination. The test of step 60 is useful in recognising an all scalar operation when mixed vector and scalar instructions are being used. It will also be noticed that when operating in a mixed vector and scalar mode, if the first operand is a scalar, then it may be any of S0 to S7, whilst if the first operand is a vector, then it may be any of S8 to S31. Providing three times the number of registers to be available within the register bank for the first operand being a vector is an adaptation to the generally greater number of registers needed to hold sequences of data values when using vector operations.

It will be appreciated that a common operation one may wish to perform is a graphics transformation. In the general case, the transformation to be performed may

be represented by a 4×4 matrix. The operand reuse in such calculations means that it is desirable that the matrix values be stored in registers that may be manipulated as vectors. In the same way, an input pixel value is usually stored in 4 registers that again should be able to be manipulated as a vector to aid reuse. The output of the
5 matrix operation will usually be scalars (accumulating the separate vector line multiplies) stored in 4 registers. If it is desired to double pump the input and output values, then you arrive at a requirement for 24 ($16 + 4 + 4$) vector registers and 8 ($4 + 4$) scalar registers.

Figure 4 is a flow diagram corresponding to that of Figure 3, but in this case
10 illustrating double precision mode. As previously mentioned, in double precision mode the register slots within the register bank 38 act as pairs to store 16 64-bit data values in logical registers D0 to D15. In this case, the encoding of the vector/scalar nature of the registers is modified from that of Figure 3 in that the tests of steps 60 and 66 now become "Is the destination one of D0 to D3?" and "Is the first operand
15 one of D0 to D3?" at steps 72 and 74 respectively.

Whilst encoding the vector/scalar nature of the registers within the register specifying fields as described above provides a significant saving in instruction bit space, it does cause some difficulties for non-commutative operations such as subtract and division. Given the register configuration $V = V \text{ op } S$, the lack of symmetry
20 between the first and second operands for non-commutative operations can be overcome without additional instructions swapping register values by extending the instruction set to include pairs of opcodes such as SUB, RSUB and DIV, RDIV that represent the two different operand options for non-commutative operations.

Figure 5 illustrates the wrapping of vectors within subsets of the register bank
25 38. In particular, in single precision mode the register bank is split into 4 ranges of registers with addresses S0 to S7, S8 to S15, S16 to S23 and S24 to S31. These ranges are disjoint and contiguous. Referring to Figure 2, the wrapping function for these subsets containing eight registers may be provided by employing 3-bit incrementers (adders) 52 within the vector control unit 50. In this way, when a subset

boundary is crossed, the incrementers will wrap back. This simple implementation is facilitated by the alignment of the subsets on eight word boundaries within the register address space.

Returning to Figure 5, a number of vector operations are illustrated to assist understanding of the wrapping of the registers. The first vector operation specifies a start register S2, a vector length of 4 (indicated by a length value within the vector control register 46 of 3) and a stride of one (indicated by a stride value within the vector control register 46 of zero). Accordingly, when an instruction is executed that is decoded to refer to register S2 as a vector with these global vector control parameters set, then the instruction will be executed 4 times respectively using the data values within the registers S2, S3, S4 and S5. As this vector does not cross a subset boundary, there is no vector wrapping.

In the second example, the starting register is S14, the length is 6 and the stride is one. This will result in the instruction being executed 6 times starting with register S14. The next register used will be S15. When the register increments by the stride again, then instead of the register used being S16, it will wrap to be register S8. The instruction is then executed further 3 times to complete the full sequence of S14, S15, S8, S9, S10 and S11.

The final example of Figure 5 shows a starting register of S25, a length of 8 and a stride of 2. The first register used will be S25 and this will be followed by S27, S29 and S31 in accordance with the stride value of 2. Following the use of register S31, the next register value will wrap back to the start of the subset, pass over register S24 in view of the stride of 2, and execute the operation using register S25. The incrementers 52 can take the form of 3-bit adders that add the stride to the current value when moving between vector registers. Accordingly, the stride can be adjusted by supplying a different stride value to the adder.

Figure 6 illustrates the wrapping of the register bank 38 within double precision mode. In this mode, the subsets of registers comprises D0 to D3, D4 to D7, D8 to D11 and D12 to D15. The minimum value input to the adder serving as the

incrementer 52 in double precision mode will be 2 - corresponding to a double precision stride of one. A double precision stride of two will require an input of 4 to the adder. The first example illustrated in Figure 6 has a start register of D0, a length of 4 and a stride one. This will result in a vector register sequence of D0, D1, D2 and
 5 D3. As no subset boundaries are crossed, there is no wrapping in this example. In the second example, the start register is D15, the length is 2 and the stride is 2. This results in a vector register sequence of D15 and D13.

Referring to Figure 2, it will be noted the load store control unit 42 has a 5-bit incrementer at its output and that load/store multiple operations are not subject to the
 10 register wrapping applied to vector operations. This enables a single load/store multiple instruction to access as many consecutive registers as it requires.

An example of an operation that makes good use of this wrapping arrangement is an FIR filter split into units of 4 signal values and 4 taps. If the syntax R8-R11 op R16-R19 represents the vector operations R8opR16, R9opR17, R10opR18 and
 15 R11opR19, then the FIR filter operation may be performed as:

Load 8 taps in R8-R15 and 8 signal values into R16-R23

R8-R11opR16-R19 and put results into R24-R27

20 R9-R12opR16-R19 and accumulate the results into R24-R27

R10-R13opR16-R19 and accumulate the results into R24-R27

R11-R14opR16-R19 and accumulate the results into R24-R27

Reload R8-R11 with new taps

25

R12-R15opR16-R19 and accumulate the results into R24-R27

R13-R8opR16-R19 and accumulate the results into R24-R27 (R15 -> R8
 wrap)

R14-R9opR16-R19 and accumulate the results into R24-R27 (R15 -> R8

wrap)

R15-R10opR16-R19 and accumulate the results into R24-R27 (R15 -> R8

wrap)

5 Reload R12 to R15 with new taps

When out of taps, reload R16-R19 with new data

R12-R15opR20-R23 and put results in R28-R31

10 R13-R8opR20-R23 and accumulate results into R28-R31 (R15 -> R8 wrap)

R14-R9opR20-R23 and accumulate results into R28-R31 (R15 -> R8 wrap)

R15-R10opR20-R23 and accumulate results into R28-R31 (R15 -> R8 wrap)

The rest as above.

15

It should be noted from the above that the loads are to different registers from the multiple accumulates and so can take place in parallel (i.e. achieves double buffering).

20 Figure 7A schematically illustrates how the main processor 24 views a coprocessor instruction. The main processor uses a bit combination of a field 76 (which may be split) within the instruction to identify the instruction as a coprocessor instruction. Within standard ARM processor instruction set, a coprocessor instruction includes a coprocessor number field 78 that the coprocessor(s) attached to the main
25 processor use to identify if a particular coprocessor instruction is targeted at them. Different types of coprocessor, such as a DSP coprocessor (e.g. the Piccolo coprocessor produced by ARM) or a floating point unit coprocessor, can be allocated different coprocessor numbers and so separately addressed within a single system using the same coprocessor bus 36. The coprocessor instructions also include an

opcode that is used by the coprocessor and three 5-bit fields respectively specifying the destination, first operand and second operand from among the coprocessor registers. In some instructions, such as a coprocessor load or store, the main processor at least partially decodes the coprocessor instruction such that the
5 coprocessor and main processor can together complete the desired data processing operation. The main processor may also be responsive to the data type encoded within the coprocessor number as part of the instruction decode it performs in such circumstances.

Figure 7B illustrates how a coprocessor supporting both double and single
10 precision operations interprets a coprocessor instruction it receives. Such a coprocessor is allocated two adjacent coprocessor numbers and uses the most significant 3 bits of the coprocessor number to identify whether it is the target coprocessor. In this way, the least significant bit of the coprocessor number is redundant for the purpose of identifying the target coprocessor and can instead be
15 used to specify the data type to be used in executing that coprocessor instruction. In this example, the data type corresponds to the data size being either single or double precision.

It can be noted that whilst in double precision mode, the number of registers is effectively reduced from 32 to 16. Accordingly, it would be possible to decrease the
20 register field size, but in that case the decode of which register to use would not be available directly from a self-contained field in a known position within the coprocessor instruction and would be dependent upon the decoding of other portions of the coprocessor instruction. This would disadvantageously complicate and possibly slow the operation of the coprocessor. Using the least significant bit of the
25 coprocessor number to encode the data type means that the opcode can be completely independent of data type which also simplifies and speeds its decode.

Figure 7C illustrates how a coprocessor supporting only a single data type that is a subset of the data types supported by the Figure 7B coprocessor interpreters the coprocessor instructions. In this case, the full coprocessor number is used to

determine whether or not to accept the instruction. In this way, if a coprocessor instruction is of a data type not supported, then it corresponds to a different coprocessor number and will not be accepted. The main processor 24 can then fall back on undefined instruction exception handling to emulate the operation on the unsupported data type.

Figure 8 illustrates a data processing system comprising an ARM core 80 serving as a main processor and communicating via a coprocessor bus 82 with a coprocessor 84 that supports both single and double precision data type. The coprocessor instruction, including the coprocessor number, is issued from the ARM core 80 on the coprocessor bus 82 when it is encountered within the instruction stream. The coprocessor 84 then compares the coprocessor number with its own numbers and if a match occurs issues an accept signal back to the ARM core 80. If no accept signal is received, then the ARM core recognises an undefined instruction exception and refers to exception handling code stored in the memory system 86.

Figure 9 illustrates the system of Figure 8 modified by replacing the coprocessor 84 with a coprocessor 88 that supports only single precision operations. In this case the coprocessor 88 recognises only a single coprocessor number. Accordingly, double precision coprocessor instructions within the original instruction stream that would be executed by the coprocessor 84 of Figure 8 are not accepted by the single precision coprocessor 88. Thus, if it is desired to execute the same code, then the undefined exception handling code within the memory system 86 can include a double precision emulation routine.

It will be noted that whilst the need to emulate double precision instructions will make the execution of these instructions slow, the single precision coprocessor 88 can be smaller and less expensive than the double precision equivalent 84 and a net benefit gained if double position instructions are sufficiently rare.

Figure 10 illustrates the instruction latch circuit within the coprocessor 84 that supports both single and double precision instructions and has two adjacent coprocessor numbers. In this case, the most significant 3 bits CP#[3:1] of the

coprocessor number within the coprocessor instruction are compared with those allocated for that coprocessor 84. In this example, if the coprocessor 84 has coprocessor numbers 10 and 11, then this comparison can be achieved by matching the most significant the bits of the coprocessor number CP#[3:1] against binary 101.

- 5 If a match occurs, then an accept signal is returned to the ARM core 80 and the coprocessor instruction is latched for execution.

Figure 11 illustrates the equivalent circuit within the single precision coprocessor 88 of Figure 9. In this case only a single coprocessor number will be recognised and single precision operations used by default. The comparison made in
10 determining whether to accept and latch the coprocessor instruction is between the full 4 bits of the coprocessor number CP#[3:0] and the single embedded coprocessor number of binary 1010.

Figure 12 is a flow illustrating how the undefined exception handling routine of the Figure 9 embodiment may be triggered to run the double precision emulation
15 code. This is achieved by detecting (step 90) if the instruction that gave rise to the undefined instruction exception is a coprocessor instruction with a coprocessor number of binary 1011. If yes, then this was intended as a double precision instruction and so can be emulated at step 92 before returning to the main program flow. Other exception types may be detected and handled by further steps if not
20 trapped by step 90.

Figure 13 illustrates the use of a format register, FPREG 200, to store information identifying the type of data stored in each 32-bit register, or data slot, of the register bank 220. As mentioned earlier, each data slot can operate individually as a single precision register for storing a 32-bit data value (a data word), or can be paired
25 with another data slot to provide a double precision register for storing a 64-bit data value (2 data words). In accordance with preferred embodiments of the present invention, the FPREG register 200 is arranged to identify whether any particular data slot has single precision or double precision data stored therein.

As illustrated in Figure 13, the 32 data slots in the register bank 220 are arranged

to provide 16 pairs of data slots. If a first data slot has a single precision data value stored therein, then in preferred embodiments the other data slot in that pair will be arranged to only store a single precision data value, and will not be linked with any other data slot in order to store a double precision data value. This ensures that any particular pair of data slots is arranged to store either two single precision data values, or one double precision data value. This information can be identified by a single bit of information associated with each pair of data slots in the register bank 220, and hence in preferred embodiments the FPREG register 200 is arranged to store 16 bits of information to identify the type of data stored in each pair of data slots of the register bank 220. It will be appreciated that the register FPREG 200 can hence be embodied as a 16-bit register, or, for consistency with other registers within the FPU coprocessor 26, can be embodied as a 32-bit register having 16 spare bits of information.

Figure 15 illustrates six pairs of data slots within the register bank 220, which can in accordance with preferred embodiments be used to store six double precision data values or twelve single precision data values. An example of data which may be stored within these data slots is shown in Figure 15, DH representing the 32 most significant bits of a double precision data value, DL indicating the 32 lowest significant bits of a double precision data value, and S representing a single precision data value.

The corresponding entries within the FPREG register 200 in accordance with preferred embodiments of the present invention are also illustrated in Figure 15. In accordance with the preferred embodiment, the value "1" is stored in the FPREG register 200 to indicate that the associated pair of data slots contains a double precision data value, and the value "0" is used to indicate that at least one of the corresponding pair of data slots contains a single precision data value, or that both data slots are uninitialised. Hence, if both data slots are uninitialised, if one of the data slots is uninitialised and the other data slot in the pair contains a single precision data value, or if both data slots in the pair contain a single precision data value, then a logic "0" value will be stored in the corresponding bit of the FPREG register 200.

As mentioned earlier, the FPU coprocessor 26 of preferred embodiments may be

used to process either single precision or double precision data values, and coprocessor instructions issued by the main processor 24 will identify whether any particular instruction is a single precision or a double precision instruction (see Figure 7B and associated description). If an instruction is accepted by the coprocessor, it will be
5 passed to the register control and instruction issue unit 48 for decoding and execution. If the instruction is a load instruction, the register control and instruction issue logic 48 will instruct the load store control unit 42 to retrieve the identified data from memory, and to store that data in the specified data slots of the register bank 220. At this stage, the coprocessor will know whether single precision or double precision data values are
10 being retrieved, and the load store control unit 42 will act accordingly. Hence, the load store control logic 42 will either pass 32-bit single precision data values, or 64-bit double precision data values, over path 225 to the register bank input logic 230 for storing in the register bank 220.

In addition to the data being loaded by the load store control unit 42 into the
15 register bank 220, data is also provided to the format register FPREG 200 to enable the necessary bits of information to be added to identify whether each pair of data slots receiving data is storing single precision or double precision data. In preferred embodiments, this data is stored in the format register FPREG 200 before data is loaded into the register bank, so that this information is available to the register bank input logic
20 230.

In preferred embodiments, the internal format of the data in the register bank 220 is the same as the external format, and hence single precision data values are stored as 32-bit data values, and double precision data values are stored as 64-bit data values within the register bank 220. Since the register bank input logic 230 has access to the
25 FPREG format register 200, it knows whether the data it is receiving is single or double precision, and so, in such an embodiment, the register bank input logic 230 merely arranges the data received over path 225 for storing in the appropriate data slot(s) of the register bank 220. However, if in alternative embodiments, the internal representation within the register bank is different to the external format, then the register bank input

logic 230 would be arranged to perform the necessary conversion. For example, a number is typically represented as 1.abc... multiplied by a base value raised to the power of an exponent. For the sake of efficiency, typical single and double precision representations do not use a data bit to represent the 1 to the left of the decimal point, but rather the 1 is taken as implied. If, for any reason, the internal representation used within the register bank 220 required the 1 to be represented explicitly, then the register bank input logic 230 would perform the necessary conversion of the data. In such embodiments, the data slots would typically be somewhat bigger than 32 bits in order to accommodate the additional data generated by the register bank input logic 230.

10 In addition to loading data values into the register bank 220, the load store control unit 42 may also load data into one or more system registers of the coprocessor 26, for example a user status and control register FPSCR 210. In preferred embodiments, the FPSCR register 210 contains user accessible configuration bits and exception status bits, and is discussed in more detail in the architectural description of the floating point unit provided at the end of the preferred embodiment description.

15 If the register control and instruction issue unit 48 receives a store instruction identifying particular data slots in the register bank 220 whose contents are to be stored to memory, then the load store control unit 42 is instructed accordingly, and the necessary data words are read out from the register bank 220 to the load store control unit 42 via the register bank output logic 240. The register bank output logic 240 has access to the FPREG register 200 contents in order to determine whether the data being read out is single or double precision data. It then applies appropriate data conversion to reverse any data conversion applied by the register bank input logic 230, and provides the data to the load store control logic 42 over path 235.

25 In accordance with the preferred embodiments of the present invention, if the store instruction is a double precision instruction, then the coprocessor 26 can be considered to be operating in a second mode of operation where instructions are applied to double precision data values. Since double precision data values contain an even number of data words, then any store instruction issued in the second mode of operation

would typically identify an even number of data slots whose contents are to be stored to memory. However, in accordance with preferred embodiments of the present invention, if an odd number of data slots are specified, then the load store control unit 42 is arranged to read the contents of FPREG register 200 and to first store those contents to
5 memory prior to storing the identified even number of data slots from the register bank 220. Typically the data slots to be transferred are identified by a base address identifying a particular data slot in the register bank, followed by a number indicating the number of data slots (i.e. number of data words), counting from the identified data slot, that are to be stored.

10 Hence, if as an example, the store instruction gives as a base address the first data slot in the register bank 220, and specifies 33 data slots, this will cause the contents of all 32 data slots to be stored to memory, but, since the specified number of data slots is odd, it will also cause the contents of the FPREG register 200 to be stored to memory.

By this approach, a single instruction can be used to store both the contents of
15 the register bank to memory, and the contents of the FPREG register 200 identifying the data types stored within the various data slots of the register bank 220. This avoids a separate instruction having to be issued to explicitly store the contents of the FPREG register 200, and hence does not so adversely affect the processing speed during a store to memory or a load from memory process.

20 In further embodiments of the present invention, this technique can be taken one stage further to enable additional system registers, such as the FPSCR register 210, to also be stored to memory, if required, using a single instruction. Hence, considering the example of a register bank 220 having 32 data slots, then, as discussed earlier, if 33 data slots are identified in the store instruction, then the FPREG register 200 will be stored to
25 memory in addition to the contents of the 32 data slots in the register bank 220.

However, if a different odd number exceeding the number of data slots in the register bank is identified, for example 35, then this can be interpreted by the load store control unit 42 as a requirement to also store the contents of the FPSCR register 210 to memory in addition to the contents of FPREG register 200 and the data slots of the register bank

220. The coprocessor may also include further system registers, for example exception registers identifying exceptions that have occurred during processing of instructions by the coprocessor. If a different odd number is identified in a store instruction, for example 37, then this can be interpreted by the load store control unit 42 as a requirement to additionally store the contents of the one or more exception registers in addition to the contents of the FPSCR register 210 the FPREG register 200, and the register bank 220.

This technique is particularly useful when the code initiating the store or load instruction is not aware of the register bank content, and the register bank content is only temporarily stored to memory for subsequent retrieval into the register bank. If the code were aware of the register bank content, then it may not be necessary for the contents of FPREG register 200 to also be stored to memory. Typical examples of code which may be unaware of the register bank content are context switch code and procedure call entry and exit routines.

In such cases, the contents of the FPREG register 200 can be efficiently stored to memory in addition to the contents of the register bank, and indeed, as discussed above, certain other system registers can also be stored as required.

Upon receipt of a subsequent load instruction, a similar process is employed. Hence, the load store control unit 42, upon receiving a double precision load instruction specifying an odd number of data slots, will be arranged to cause the contents of FPREG register 200 to be loaded into the FPREG register 200, followed by the contents of any system registers indicated by the number of slots identified in the load instruction, followed by an even number of data words to be stored in the specified data slots of the register bank 220. Hence, considering the earlier discussed example, if the number of data slots specified in the load instruction is 33, then the FPREG register contents will be loaded into the FPREG register 200, followed by the contents of the 32 data slots. Similarly, if the number of data slots specified in the load instruction is 35, then the contents of the FPSCR register 210 will also be loaded into the FPSCR register in addition to the above mentioned contents. Finally, if the number of data slots specified

is 37, then the contents of any exception registers will also be loaded into those exception registers in addition to the above mentioned contents. Clearly, it will be appreciated by those skilled in the art that the particular actions associated with particular odd numbers is entirely arbitrary, and can be varied as desired.

5 Figure 14 is a flow diagram illustrating operation of the register control and instruction issue logic 48 in accordance with preferred embodiments of the present invention when executing store and load instructions. Firstly, at step 300, the number of data words (which is identical to the number of data slots in preferred embodiments) is read from the instruction, along with the first register number, i.e. the base register,
10 identified in the instruction. Then, at step 310, it is determined whether the instruction is a double precision instruction, as mentioned previously this information being available to the coprocessor at this stage since the instruction identifies whether it is a double precision or a single precision instruction.

 If the instruction is a double precision instruction, then the process proceeds to
15 step 320, where it is determined whether the number of words specified in the instruction is odd. Assuming for the sake of this embodiment that the technique is not used to selectively transfer various system registers in addition to the FPREG register 200, then if the number of words is odd, this will indicate that the contents of the FPREG register 200 should be transferred, and accordingly at step 325, the contents of
20 the FPREG register are transferred by the load store control unit 42. Then, the number of words is decremented by 1 at step 327, and the process proceeds to step 330. If, at step 320, the number of words was determined to be even, then the process proceeds directly to step 330.

 At step 330, it is determined whether the number of words is greater than zero.
25 If not, then the instruction is deemed completed, and the process exits at step 340. However, if the number of words is greater than zero, then the process proceeds to step 332, where a double precision data value (i.e. the contents of two data slots) is transferred to or from the first specified register number. Then, at step 334, the number of words is decremented by 2, and at step 336, the register number is incremented by 1.

As discussed earlier, for a double precision instruction, a register actually consists of two data slots, and hence incrementing the register count by one is equivalent to incrementing the data slot number by 2.

Then the procedure returns to step 330, where it is determined whether the
5 number of words is still greater than zero, and if so the process is repeated. When the number of words reaches zero, then the process is exited at step 340.

If at step 310, it was determined that the instruction was not a double precision instruction, then the process proceeds to step 350, where it is again determined whether the number of words is greater than zero. If so, the process proceeds to step 352, where
10 a single precision data value is transferred to or from the first register number identified in the instruction. Then, at step 354, the number of words is decremented by one, and at step 356 the register number count is incremented by one so as to point at the next data slot. Then the process returns to step 350, where it is determined whether the number of words is still greater than zero. If so, the process is repeated, until such time as the
15 number of words is equal to zero, at which time the process is exited at step 360.

The above approach provides a great deal of flexibility when executing code which is unaware of the register bank contents, for example context switch code or procedure call entry and exit sequences. In these cases, the operating system is not aware of the contents of the registers, and it is desirable not to have to treat the registers
20 differently, dependent on their contents. The above approach allows these code routines to be written with a single store or load instruction specifying an odd number of data words. If the coprocessor requires the use of the register content information, it will interpret the odd number of data words in the instruction as a requirement to also store to memory or load from memory the format information required to identify the
25 contents of the data in the register bank. This flexibility removes the need for unique operating system software to support coprocessors that require the register content information.

This technique also removes the necessity for loading and storing the register content information in a separate operation within the code. Since the option to load and

store the register content information is incorporated in the instruction, no additional memory access is required. This reduces the code length and potentially saves time.

An architectural description of a floating point unit incorporating the above described techniques is given below:

1. Introduction

- 5 The VFPv1 is a floating point system (FPS) architecture designed to be implemented as a coprocessor for use with ARM processor modules. Implementations of this architecture may incorporate features in either hardware or software, or an implementation may use software to complement the functionality or provide IEEE 754 compatibility. This specification intends to achieve full IEEE 754 compatibility
- 10 using a combination of hardware and software support.

Two coprocessor numbers are used by VFPv1; 10 is used for operations with single precision operands, while 11 is used for operations with double precision operands. Conversion between single and double precision data is accomplished with 2

15 conversion instructions which operate in the source operand coprocessor space.

Features of the VFPv1 architecture include:

- Full compatibility with IEEE 754 in hardware with support code.
- 20 • 32 single precision registers, each addressable as a source operand or a destination register.
- 16 double precision registers, each addressable as a source operand or a destination register. (Double precision registers overlap physical single precision registers)
- 25 • Vector mode provides for a significant increase in floating point code density and concurrency with load and store operations.
- 4 banks of 8 circulating single precision registers or 4 banks of 4 circulating double precision registers to enhance dsp and graphics operations.
- Denormal handling option selects between IEEE 754 compatibility (with intended

support from the floating point emulation package) or fast flush-to-zero capability.

- Intended for implementation with a fully pipelined chained multiply-accumulate with IEEE 754 compatible results.
- Fast floating point to integer conversion for C, C++, and Java with the FFTOSIZ
5 instruction.

Implementers may choose to implement the VFPv1 completely in hardware or utilize a combination of hardware and support code. The VFPv1 may be implemented completely in software.

10 2. Terminology

This specification uses the following terminology:

- Automatic exception** - An exceptional condition which will always bounce to the
15 support code regardless of the value of the respective exception enable bit. The choice of which, if any, exceptions are Automatic is an implementation option. See Section 0,

20 6. Exception Processing.

- Bounce** - An exception reported to the operating system which will be handled by the support code entirely without calling user trap handlers or otherwise interrupting the normal flow of user code.

- 25 **CDP** - 'Coprocessor Data Processing' For the FPS, CDP operations are arithmetic operations rather than load or store operations.

ConvertToUnsignedInteger(Fm) - Conversion of the contents in Fm to a unsigned 32-bit integer value. The result is dependent on the rounding mode for final rounding

and handling of floating point values outside the range of a 32-bit unsigned integer. The INVALID exception is possible if the floating point input value is negative or too large for a 32-bit unsigned integer.

- 5 **ConvertToSignedInteger(Fm)** - Conversion of the contents in Fm to a signed 32-bit integer value. The result is dependent on the rounding mode for final rounding and handling of floating point values outside the range of a 32-bit signed integer. The INVALID exception is possible if the floating point input value is too large for a 32-bit signed integer.

10

ConvertUnsignedIntToSingle/Double(Rd) - Conversion of the contents of an ARM register (Rd), interpreted as a 32-bit unsigned integer, to a single or double precision floating point value. If the destination precision is single, an INEXACT exception is possible in the conversion operation.

15

ConvertSignedIntToSingle/Double(Rd) - Conversion of the contents of an ARM register (Rd), interpreted as a 32-bit signed integer, to a single or double precision floating point value. If the destination precision is single, an INEXACT exception is possible in the conversion operation.

20

Denormalized value - A representation of a value in the range $(-2^{E_{min}} < x < 2^{E_{min}})$. In the IEEE 754 format for single and double precision operands, a denormalized value, or denormal, has a zero exponent and the leading significand bit is 0 rather than 1.

- The IEEE 754-1985 specification requires that the generation and manipulation of
25 denormalized operands be performed with the same precision as with normal operands.

Disabled exception - An exception which has its associated Exception Enable bit in the FPCSR set to 0 is referred to as 'disabled.' For these exceptions the IEEE 754

specification defines the correct result to be returned. An operation which generates an exception condition may bounce to the support code to produce the IEEE 754 defined result. The exception will not be reported to the user exception handler

- 5 **Enabled exception** - An exception with the respective exception enable bit set to 1. In the event of an occurrence of this exception a trap to the user handler will be taken. An operation which generates an exception condition may bounce to the support code to produce the IEEE 754 defined result. The exception will then be reported to the user exception handler.

10

Exponent - The component of a floating point number that normally signifies the integer power to which two is raised in determining the value of the represented number. Occasionally the exponent is called the signed or unbiased exponent.

- 15 **Fraction** - The field of the significand that lies to the right of its implied binary point.

Flush-To-Zero Mode - In this mode all values in the range $(-2^{E_{min}} < x < 2^{E_{min}})$ after rounding are treated as zero, rather than converted to a denormalized value.

- 20 **High(Fn/Fm)** - The upper 32 bits [63:32] of a double precision value as represented in memory.

- IEEE 754-1985** - "IEEE Standard for Binary Floating-Point Arithmetic", ANSI/IEEE Std 754-1985, The Institute of Electrical and Electronics Engineers, Inc. New York, New York, 10017. The standard, often referred to as the IEEE 754 standard, which defines data types, correct operation, exception types and handling, and error bounds for floating point systems. Most processors are built in compliance with the standard in either hardware or a combination of hardware and software.

Infinity - An IEEE 754 special format used to represent ∞ . The exponent will be maximum for the precision and the significand will be all zeros.

Input exception - An exception condition in which one or more of the operands for a given operation are not supported by the hardware. The operation will bounce to support code for completion of the operation.

Intermediate result - An internal format used to store the result of a calculation before rounding. This format may have a larger exponent field and significand field than the destination format.

Low(Fn/Fm) - The lower 32 bits [31:0] of a double precision value as represented in memory.

- 15 **MCR** - "Move to Coprocessor from ARM Register" For the FPS this includes instructions which transfer data or control registers between an ARM register and a FPS register. Only 32 bits of information may be transferred using a single MCR class instruction.
- 20 **MRC** - "Move to ARM Register from Coprocessor" For the FPS this includes instructions which transfer data or control registers between the FPS and an ARM register. Only 32 bits of information may be transferred using a single MRC class instruction.
- 25 **NaN** - Not a number, a symbolic entity encoded in a floating point format. There are two types of NaNs, signalling and non-signalling, or quiet. Signalling NaNs will cause an Invalid Operand exception if used as an operand. Quiet NaNs propagate through almost every arithmetic operation without signalling exceptions. The format for a NaN has the exponent field of all 1's with the significand non-zero. To represent

a signalling NaN the most significant bit of the fraction is zero, while a quiet NaN will have the bit set to a one.

- Reserved** - A field in a control register or instruction format is 'reserved' if the field is to be defined by the implementation or would produce UNPREDICTABLE results if the contents of the field were not zero. These fields are reserved for use in future extensions of the architecture or are implementation specific. All Reserved bits not used by the implementation must be written as zero and will be read as zero.
- 5
- 10 **Rounding Mode** - The IEEE 754 specification requires all calculations to be performed as if to an infinite precision, that is, a multiply of two single precision values must calculate accurately the significand to twice the number of bits of the significand. To represent this value in the destination precision rounding of the significand is often required. The IEEE 754 standard specifies four rounding modes -
- 15 round to nearest (RN), round to zero, or chop (RZ), round to plus infinity (RP), and round to minus infinity (RM). The first is accomplished by rounding at the half way point, with the tie case rounding up if it would zero the lsb of the significand, making it 'even.' The second effectively chops any bits to the right of the significand, always rounding down, and is used by the C, C++, and Java languages in integer conversions.
- 20 The later two modes are used in interval arithmetic.

Significand - The component of a binary floating point number that consists of an explicit or implicit leading bit to the left of its implied binary point and a fraction field to the right.

25

Support Code - Software which must be used to complement the hardware to provide compatibility with the IEEE 754 standard. The support code is intended to have two components: a library of routines which perform operations beyond the scope of the hardware, such as transcendental computations, as well as supported functions, such

as divide with unsupported inputs or inputs which may generate an exception; and a set of exception handlers which process exceptional conditions in order to provide IEEE 754 compliance. The support code is required to perform implemented functions in order to emulate proper handling of any unsupported data type or data representation (e.g., denormal values or decimal datatypes). The routines may be written to utilize the FPS in their intermediate calculations if care is taken to restore the users' state at the exit of the routine.

Trap - An exceptional condition which has the respective exception enable bit set in the FPSCR. The user's trap handler will be executed.

UNDEFINED - Indicates an instruction that generates an undefined instruction trap. See the ARM Architectural Reference Manual for more information on ARM exceptions.

15

UNPREDICTABLE - The result of an instruction or control register field value that cannot be relied upon. UNPREDICTABLE instructions or results must not represent security holes, or halt or hang the processor, or any parts of the system.

Unsupported Data - Specific data values which are not processed by the hardware but bounced to the support code for completion. These data may include infinities, NaNs, denormal values, and zeros. An implementation is free to select which of these values will be supported in hardware fully or partially, or will require assistance from support code to complete the operation. Any exception resulting from processing unsupported data will be trapped to user code if the corresponding exception enable bit for the exception is set.

3. Register File

3.1 Introductory Notes

The architecture provides 32 single precision and 16 double precision registers, all individually addressable within a fully defined 5-bit register index as source or destination operands.

5

The 32 single precision registers are overlapped with the 16 double precision registers, i.e., a write of a double precision data to D5 will overwrite the contents of S10 and S11. It is the job of the compiler or the assembly language programmer to be aware of register usage conflicts between the use of a register as a single precision data storage and as half of a double precision data storage in an overlapped implementation. No hardware is provided to insure register use is limited to one precision, and the result is UNPREDICTABLE if this is violated.

10

VFPv1 provides access to these registers in a scalar mode, in which one, two, or three operand registers are used to produce a result which is written into a destination register, or in vector mode, in which the operands specified refer to a group of registers. VFPv1 supports vector operations for up to eight elements in a single instruction for single precision operands and up to 4 elements for double precision operands.

15

20

Table 1. LEN Bit Encodings

<u>LEN</u>	<u>Vector Length Encoding</u>
000	Scalar
001	Vector length 2
010	Vector length 3
011	Vector length 4
100	Vector length 5

<u>LEN</u>	<u>Vector Length Encoding</u>
101	Vector length 6
110	Vector length 7
111	Vector length 8

Vector mode is enabled by writing a non-zero value to the LEN field. If the LEN field contains 0, the FPS operates in scalar mode, and the register fields are interpreted as addressing 32 individual single precision registers or 16 double precision registers in a flat register model. If the LEN field is non-zero, the FPS operates in vector mode, and the register fields are as addressing vectors of registers. See Table 1 for encoding of the LEN field.

- 10 A means of mixing scalar and vector operations without changing the LEN field is available through the specification of the destination register. Scalar operations may be specified while in vector mode if the destination register is in the first bank of registers (S0 - S7 or D0 - D3). See Section 0 for more information:

3.2 Single Precision Register Usage

- 15 If the LEN field in the FPSCR is 0, 32 single precision registers are available numbered S0 through S31. Any of the registers may be used as a source or destination register.

31	31	31	31	0
S0	S8	S16	S24	
S1	S9	S17	S25	
S2	S10	S18	S26	
S3	S11	S19	S27	
S4	S12	S20	S28	
S5	S13	S21	S29	
S6	S14	S22	S30	
S7	S15	S23	S31	

Illustration 1. Single Precision Register

Map

The single precision (coprocessor 10) register map may be drawn as shown in Illustration 1.

- 5 If the LEN field in the FPSCR is greater than 0, the register file behaves as 4 banks of 8 circulating registers, as shown in Illustration 2. The first bank of vector registers, V0 through V7, overlap with scalar registers S0 through S7, and are addressed as scalars or vectors according to the registers selected for each operand. See Section 0, 3.4 Register Usage, for more information.

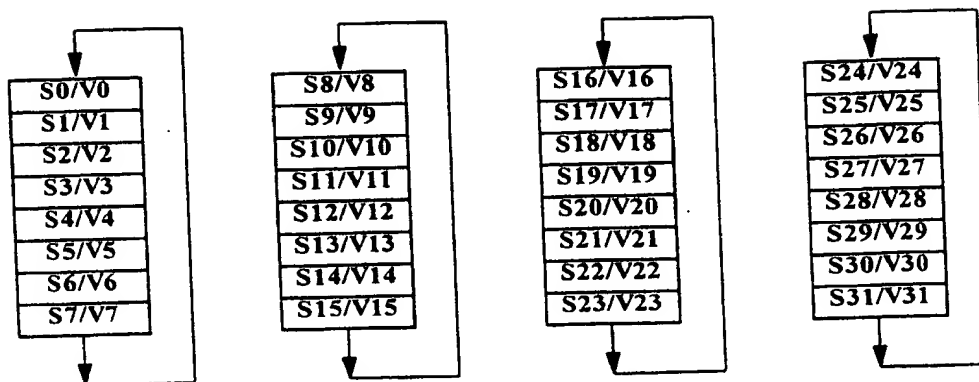


Illustration 2. Circulating Single Precision Registers

For example, if the LEN in the FPSCR is set to 3, referencing vector V10 will cause registers S10, S11, S12, and S13 to be involved in a vector operation. Similarly, V22 would involve S22, S23, S16, and S17 in the operation. When the register file is accessed in vector mode, the register following V7 in order is V0; similarly, V8 follows V15, V16 follows V23, and V24 follows V31.

3.3 Double Precision Register Usage

If the LEN field in the FPSCR is 0, 16 double precision scalar registers are available.

63	0 63	0
D0	D8	
D1	D9	
D2	D10	
D3	D11	
D4	D12	
D5	D13	
D6	D14	
D7	D15	

Illustration 3. Double Precision
Register Map

Any of the registers may be used as a source or destination register. The register map may be drawn as shown in Illustration 3.

If the LEN field in the FPSCR is greater than 0, 4 scalar registers and 16 vector registers, in 4 banks of 4 circulating registers, are available as shown in Illustration 4.

- 5 The first bank of vector registers, V0 through V3, overlap with scalar registers D0 through D3. The registers are addressed as scalars or according to the registers selected for each operand. See Section 0, 3.4 Register Usage, for more information.

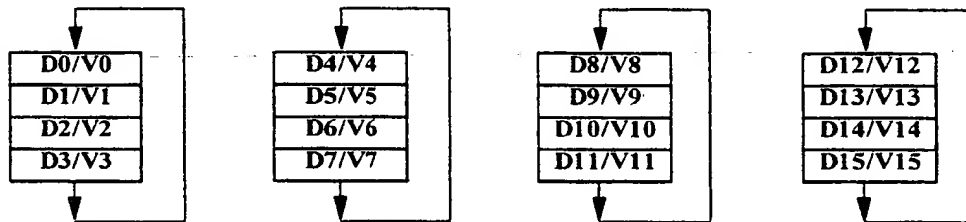


Illustration 4. Circulating Double Precision Registers

- 10 As with the single precision examples in Section 0, the double precision registers are circulating within the four banks.

3.4 Register Usage

- Three operations between scalars and vectors are supported: (OP_2 may be any of the two operand operations supported by the floating point coprocessor; OP_3 may be any of the three operand operations.)
- 15

For the following descriptions, the 'first bank' of the register file is defined as registers S0 - S7 for single precision operations and D0 - D3 for double precision operations.

20

- $ScalarD = OP_2 \text{ ScalarA}$ or $ScalarD = ScalarA \text{ } OP_3 \text{ ScalarB}$ or $ScalarD = ScalarA * ScalarB + ScalarD$
- $VectorD = OP_2 \text{ ScalarA}$ or $VectorD = ScalarA \text{ } OP_3 \text{ VectorB}$ or $VectorD =$

ScalarA * VectorB + VectorD

- VectorD = OP₂ VectorA or VectorD = VectorA OP₃ VectorB or VectorD = VectorA * VectorB + VectorD

3.4.1 Scalar Operations

5 Two conditions will cause the FPS to operate in scalar mode:

1? The LEN field in the FPSCR is 0. Destination and source registers may be any of the scalar registers, 0 through 31 for single precision operations and 0 through 15 for double precision operations. The operation will be performed only on the registers explicitly specified in the instruction.

2? The destination register is in the first bank of the register file. The source scalars may be any of the other registers. This mode allows the intermixing of scalar and vector operations without having to change the LEN field in the FPSCR.

15

3.4.2 Operations Involving a Scalar and Vector Source with a Vector Destination

To operate in this mode, the LEN field in the FPSCR is greater than zero, and the destination register is not in the first bank of the register file. The scalar source registers may be any register in the first bank of the register file while any of the remaining registers may be used for VectorB. Note that the behavior is UNPREDICTABLE if the source scalar register is a member of VectorB or if VectorD overlaps VectorB in less than LEN elements. I.e., Vector D and VectorB must be either the same vector or completely distinct in all members. See the summary tables in Section 0.

25

3.4.3 Operations Involving Only Vector Data

To operate in this mode, the LEN field in the FPSCR is greater than zero and the destination vector register is not in the first bank of the register file. The individual elements of the VectorA vector are combined with the corresponding element in

VectorB and written to VectorD. Any register not in the first bank of the register file is available for Vector A, while all vectors are available for VectorB. As in the second case, the behavior is UNPREDICTABLE if the either of the source vectors and the destination vector overlap in less than LEN elements. They must be identical or
 5 completely distinct in all members. See the summary tables in Section 0.

Note that for the FMAC family of operations the destination register or vector is always the accumulate register or vector.

3.4.4 Operation Summary Tables

- 10 The following tables present the register usage options for single and double precision 2 and 3 operand instructions. 'Any' refers to availability of all registers in the precision for the specified operand.

Table 2. Single Precision 3-Operand Register Usage

LEN field	Destination Reg	First Source Reg	Second Source Reg	Operation Type
0	Any	Any	Any	$S = S \text{ op } S \text{ or } S = S * S + S$
non-0	0-7	Any	Any	$S = S \text{ op } S \text{ or } S = S * S + S$
non-0	8-31	0-7	Any	$V = S \text{ op } V \text{ or } V = S * V + V$
non-0	8-31	8-31	Any	$V = V \text{ op } V \text{ or } V = V * V + V$

Table 3. Single Precision 2-Operand Register Usage

LEN field	Destination Reg	Source Reg	Operation Type
0	Any	Any	$S = \text{op } S$
non-0	0-7	Any	$S = \text{op } S$
non-0	8-31	0-7	$V = \text{op } S$
non-0	8-31	8-31	$V = \text{op } V$

Table 4. Double Precision 3-Operand Register Usage

LEN field	Destination Reg	First Source Reg	Second Source Reg	Operation Type
0	Any	Any	Any	$S = S \text{ op } S \text{ or } S = S * S + S$
non-0	0-3	Any	Any	$S = S \text{ op } S \text{ or } S = S * S + S$
non-0	4-15	0-3	Any	$V = S \text{ op } V \text{ or } V = S * V + V$
non-0	4-15	4-15	Any	$V = V \text{ op } V \text{ or } V = V * V + V$

Table 5. Double Precision 2-Operand Register Usage

LEN field	Destination Reg	Source Reg	Operation Type
0	Any	Any	$S = \text{op } S$
non-0	0-3	Any	$S = \text{op } S$
non-0	4-15	0-3	$V = \text{op } S$
non-0	4-15	4-15	$V = \text{op } V$

4. Instruction Set

FPS instructions may be divided into three categories:

- MCR and MRC - Transfer operations between the ARM and the FPS
- 5 • LDC and STC - Load and store operations between the FPS and memory
- CDP - Data processing operations

4.1 Instruction Concurrency

The intent of the FPS architectural specification is concurrency on two levels: pipelined functional units and parallel load/store operation with CDP functions. A
10 significant performance gain is available by supporting load and store operations which do not have register dependencies with currently processing operations to execute in parallel with these operations.

4.2 Instruction Serialization

The FPS specifies a single instruction that causes the FPS to busy-wait the ARM until
15 all currently executing instructions have completed and the exception status of each is known. If an exception is pending, the serializing instruction will be aborted and exception processing will begin in the ARM. The serializing instructions in the FPS is:

- 20 • FMOVX - read or write to a floating point system register

Any read or write to a floating point system register will be stalled until the current instructions have completed. An FMOVX to the System ID Register (FPSID) will trigger an exception caused by the preceding floating point instruction. Performing a
25 read/modify/write (using FMOVX) on the User Status and Control Register (FPSCR) can be used to clear the exception status bits (FPSCR[4:0]).

4.3 Conversion Involving Integer Data

The conversion between floating point and integer data is a two step process in the FPS made up of a data transfer instruction involving the integer data and a CDP

instruction performing the conversion. If any arithmetic operation is attempted on the integer data in the FPS register while in integer format the results are UNPREDICTABLE and any such operation should be avoided.

5 4.3.1 Conversion of integer data to floating point data in a FPS register

Integer data may be loaded into a floating point single precision register from either an ARM register, using a MCR FMOVS instruction. The integer data in the FPS register may then be converted into a single or double precision floating point value with the integer-to-float family of operations and written to a destination FPS register.

- 10 The destination register may be the source register if the integer value is no longer needed. The integer may be a signed or unsigned 32-bit quantity.

4.3.2 Conversion of floating point data in an FPS register to integer data

- 15 A value in a FPS single or double precision register may be converted to signed or unsigned 32-bit integer format with the float-to-integer family of instructions. The resulting integer is placed in the destination single precision register. The integer data may be stored to an ARM register using the MRC FMOVS instruction.

4.4 Register File Addressing

- 20 Instructions operating in single precision space ($S = 0$) will use the 5 bits available in the instruction field for operand access. The upper 4 bits are contained in the operand fields labeled Fn, Fm, or Fd; the least significant bit of the address is in N, M, or D, respectively.

- Instructions operating in double precision space ($S = 1$) will use only the upper 4 bits 25 of the operand address. These 4 bits are contained in the Fn, Fm, and Fd fields. The N, M, and D bits must contain 0 when the corresponding operand field contains an operand address.

4.5 MCR (Move to Coprocessor from ARM Register)

The MCR operations involve the transfer or use of data in ARM registers by the FPS.

This includes moving data in single precision format from an ARM register or in double precision format from a pair of ARM registers to an FPS register, loading a signed or unsigned integer value from an ARM register to a single precision FPS register, and loading a control register with the contents of an ARM register.

5

The format for an MCR instruction is given in Illustration 5.

31	28	27	24	23	21	20	19	16	15	12	11	8	7	6	5	4	3	0
COND	1	1	1	0	Opcode	0	Fn	Rd	1	0	1	S	N	R	R	1	Reserved	

Illustration 5. MCR Instruction Format

Table 6. MCR Bit Field Definitions

Bit Field	Definition
Opcode	3-bit operation code (See Table 7)
Rd	ARM Source register encoding
S	Operation operand size. 0 - Single precision operands 1 - Double precision operands
N	Single precision operations: Destination register lsb Double precision operations: Must be set to 0 or the operation is UNDEFINED System register moves Reserved

<u>Bit Field</u>	<u>Definition</u>
Fn	<p>Single precision operations: Destination register address upper 4 bits</p> <p>Double precision operations: Destination register address</p> <p>System register moves: 0000 - FPID (Coprocessor ID number) 0001 - FPSCR (User Status and Control Register) 0100 - FPREG (Register File Content Register) Other register encodings are Reserved and may be different on various implementations.</p>
R	Reserved bits

Table 7. MCR Opcode Field Definition

<u>Opcode Field</u>	<u>Name</u>	<u>Operation</u>
000	FMOVS	Fn = Rd (32 bits, coprocessor 10)
000	FMOVLD	Low(Fn) = Rd (Double precision low 32 bits, coprocessor 11)
001	FMOVHD	High(Fn) = Rd (Double precision high 32 bits, coprocessor 11)
010- 110	Reserved	
111	FMOVX	System Reg = Rd (coprocessor 10 space)

Note: Only 32-bit data operations are supported by FMOV[S, HD, LD] instructions. Only the data in the ARM register or single precision register is moved by the FMOVS operation. To transfer a double precision operand from 2 ARM registers the FMOVLD and FMOVHD instructions will move the lower half and the upper half, respectively.

4.6 MRC (Move to ARM Register from Coprocessor / Compare Floating Registers)

The MRC operations involve the transfer of data in an FPS register to an ARM register. This includes moving a single precision value or the result of a conversion of a floating point value to integer to an ARM register or a double precision FPS register to two ARM registers, and modifying the status bits of the CPSR with the results of a previous floating point compare operation.

The format of the MRC instruction is shown in Illustration 6.

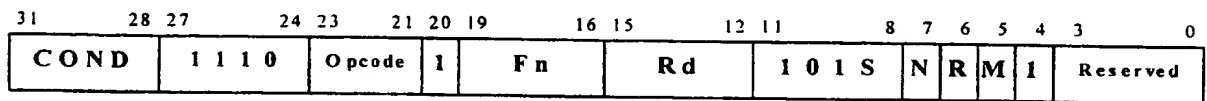


Illustration 6. MRC Instruction Format

Table 8. MRC Bit Field Definitions

Bit Field	Definition
Opcode	3-bit FPS operation code (See Table 9)
Rd	ARM destination* register encoding
S	Operation operand size. 0 - Single precision operands 1 - Double precision operands

<u>Bit Field</u>	<u>Definition</u>
N	Single precision operations: Destination register lsb Double precision operations: Must be set to 0 or operation is UNDEFINED System register moves Reserved
M	Reserved
Fn	Single precision operations: Destination register address upper 4 bits Double precision operations: Destination register address System register moves: 0000 - FPID (Coprocessor ID number) 0001 - FPSCR (User Status and Control Register) 0100 - FPREG (Register File Content Register) Other register encodings are Reserved and may be different on various implementations.
Fm	Reserved
R	Reserved

* For the FMOVX FPSCR instruction, if the Rd field contains R15 (1111), the upper 4 bits of the CPSR will be updated with the resulting condition codes.

5 Table 9. MRC Opcode Field Definition

Opcode Field	Name	Operation
000	FMOVS	Rd = Fn (32 bits, coprocessor 10)
000	FMOVLD	Rd = Low(Fn) Lower 32 bits of Dn are transferred. (Double precision low 32 bits, coprocessor 11)
001	FMOVHD	Rd = High(Fn) Upper 32 bits of Dn are transferred. (Double precision high 32 bits, coprocessor 11)
010 - 110	Reserved	
111	FMOVX	Rd = System Reg

Note: See the Note for MCR FMOV instruction.

4.7 LDC/STC (Load/Store FPS Registers)

- 5 LDC and STC operations transfer data between the FPS and memory. Floating point data may be transferred in either precision in a single data transfer or in multiple data transfers, with the ARM address register updated or left unaffected. Both full descending stack and empty ascending stack structures are supported, as well as multiple operand access to data structures in the move multiple operations. See Table 10 11 for a description of the various options for LDC and STC.

The format of the LDC and STC instructions is shown in Illustration 7.

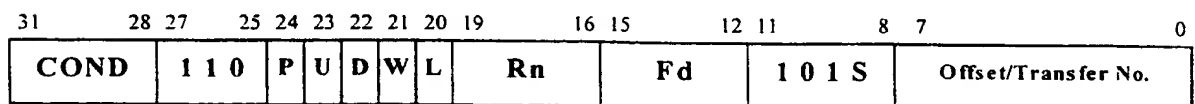


Illustration 7. LDC/STC Instruction Format

Table 10. LDC/STC Bit Field Definitions

<u>Bit Field</u>	<u>Definition</u>
P	Pre/Post Indexing (0=post, 1=pre)
U	Up/Down bit (0=down, 1=up)
D	Single precision operations: Source/Destination register lsb Double precision operations: Must be set to 0
W	Write-back bit (0=no writeback, 1=writeback)
L	Direction bit (0=store, 1=load)
Rn	ARM Base register encoding
Fd	Single precision operations: Source/Destination register address upper 4 bits Double precision operations: Source/Destination register address
S	Operation operand size. 0 - Single precision operands 1 - Double precision operands
Offset/ Transfer No.	Unsigned 8-bit offset or number of single precision (double the count of double precision registers) registers to transfer for FLDM(IA/DB) and FSTM(IA/DB). The maximum number of words in a transfer is 16, allowing for 16 single precision values or 8 double precision values.

4.7.1 General notes for load and store operations

Loading and storing multiple registers will do so linearly through the register file without wrapping across 4 or 8 register boundaries used by the vector operations. Attempting to load past the end of the register file is UNPREDICTABLE.

5

If the offset for a double load or store multiple contains an odd register count 17 or less, the implementation may write another 32-bit data item or read another 32-bit data item, but is not required to do so. This additional data item may be used to identify the contents of the registers as they are loaded or stored. This is useful in

10 implementations in which the register file format is different from the IEEE 754 format for the precision and each register has type information which is required to identify it in memory. If the offset is odd and the number is greater than the number of single precision registers, this may be used to initiate a context switch of the registers and all the system registers.

15

Table 11. Load and Store Addressing Mode Options

P	W	Offset/ <u>Transfer</u> <u>No.</u>	Addressing <u>Mode</u>	<u>Name</u>
Type 0 Transfer: Load/Store multiple with no writeback				
0	0	Number of registers to transfer	FLDM<cond><S/D> Rn, <register list> FSTM<cond><S/D> Rn, <register list>	Load/ Store Multiple

P	W	Offset/ Transfer No.	Addressing Mode	Name
<p>Load/store multiple registers from a starting address in Rn and no modification of Rn. The number of registers may be 1 to 16 for single precision, 1 to 8 for double precision. The offset field contains the number of 32-bit transfers. This mode may be used to load a transform matrix for graphics operations and a point for the transform.</p>				
<p>Examples:</p> <p>FLDMEQS r12, {f8-f11} ; loads 4 single from the address in r12 to 4 fp registers s8, s9, s10, and r12 is unchanged</p> <p>FSTMEQD r4, {f0} ; stores one double from d0 to the address in r4. r4 is unchanged.</p>				

P	W	Offset/ Transfer No.	Addressing Mode	Name
Type 1 Transfer: Load/Store multiple with post-index of Rn and writeback				
0	1	Number of registers to transfer	FLDM<cond>IA<S/D> Rn!, <register list> FSTM<cond>IA<S/D> Rn!, <register list>	Load/Store Multiple
<p>Load/Store multiple registers from a starting address in Rn and writeback of the next address after the last transfer to Rn. The offset field is the number of 32-bit transfers. The writeback to Rn is Offset*4. The maximum number of words transferred in a load multiple is 16. The U bit must be set to 1. This is used for storing into an empty ascending stack or loading from a full descending stack, or storing a transformed point and incrementing the pointer to the next point, and for loading and storing multiple data in a filter operation.</p>				
<p>Example:</p> <p>FLDMEQIAS r13!, {f12-f15} : loads 4 singles from the address in r13 to 4 fp registers s12,s13,s14, and s15, updating r13 with the address pointing to the next data the series.</p>				

P	W	Offset/ Transfer No.	Addressing Mode	Name
Type 2 Transfer: Load/Store one register with pre-index or Rn and no writeback				
1	0	Offset	FLD<cond><S/D> [Rn, #+/-offset], Fd FST<cond><S/D> [Rn, #+/-offset], Fd	Load/Store with Offset
Load/Store single register with pre-increment of the address in Rn and no writeback. The offset value is Offset*4, and is added (U=1) or subtracted (U=0) from Rn to generate the address. This is useful for operand access into a structure and is the typical method used to access memory for floating point data.				
Example: FSTEQD f4, [r8, #+8] ; Stores a double to d4 from the address in r8 offset by 32 (8 * 4) bytes. r8 is unchanged.				

P	W	Offset/ Transfer No.	Addressing Mode	Name
Type 3 Transfer: Load/Store multiple registers with pre-index and writeback				
1	1	Number of registers to transfer	FLDM<cond>DB<S/D> Rn!, <register list> FSTM<cond>DB<S/D> Rn!, <register list>	Load/Store Multiple with Pre- Decrement
<p>Load/Store multiple registers with pre-decrement of the address in Rn and writeback of the new target address to Rn. The offset field contains the number of 32-bit transfers. The writeback value is the Offset*4, subtracted (U=0) from Rn. This mode is used for storing to a full descending stack or loading from an empty ascending stack.</p> <p>Example:</p> <p>FSTMEQDBS r9!,{f27-f29} ; store 3 singles from s27, s28, and s29 to a full descending stack with the last entry address contained in r9. r9 is updated to point to the new last entry.</p>				

4.7.2 LDC/STC Operation Summary

Table 12 lists the allowable combinations for the P, W, and U bits in the LDC/STC opcode and the function of the offset field for each valid operation.

5

Table 12. LDC/STC Operation Summary

P	W	U	Offset Field	Operation
0	0	0		UNDEFINED
0	0	1	Reg Count	FLDM/FSTM

P	W	U	Offset Field	Operation
0	1	0		UNDEFINED
0	1	1	Reg Count	FLDMIA/FSTMIA
1	0	0	Offset	FLD/FST
1	0	1	Offset	FLD/FST
1	1	0	Reg Count	FLDMDB/FSTMDB
1	1	1		UNDEFINED

4.8 CDP (Coprocesor Data Processing)

CDP instructions include all data processing operations which involve operands from the floating point register file and produce a result which will be written back to the register file. Of special interest is the FMAC (multiply-accumulate chained)

- 5 operation, an operation performing a multiply on two of the operands and adding a third. This operation differs from fused multiply-accumulate operations in that an IEEE rounding operation is performed on the product before the addition of the third operand. This allows Java code to utilize the FMAC operation to speed up multiply-accumulate operations over the separate multiply then add operations.

10

- Two instructions in the CDP group are useful in conversion of a floating point value in a FPS register to its integer value. FFTOUI[S/D] performs a conversion of the contents of a single or double precision to an unsigned integer in a FPS register, using the current rounding mode in the FPSCR. FFTOSI[S/D] performs the conversion to a
- 15 signed integer. FFTOUIZ[S/D] and FFTOSIZ[S/D] perform the same functions but override the FPSCR rounding mode for the conversion and truncates any fraction bits. The functionality of FFTOSIZ[S/D] is required by C, C++, and Java in float to integer conversions. The FFTOSIZ[S/D] instructions provide this capability without requiring adjustment of the rounding mode bits in the FPSCR to RZ for the

conversion, reducing the cycle count for the conversion to only that of the FFTOSIZ[S/D] operation, saving 4 to 6 cycles.

Compare operations are performed using the CDP CMP instructions followed by a
 5 MRC FMOVX FPSCR instruction to load the ARM CPSR flag bits with the resulting FPS flag bits (FPSCR[31:28]). The compare operations are provided with and without the potential for an INVALID exception if one of the compare operands is a NaN. The FCMP and FCMP0 will not signal the INVALID if one of the compare operands is a NaN, while the FCMPE and FCMPE0 will signal the exception. The FCMP0 and
 10 FCMPE0 compare the operand in the Fm field with 0 and set the FPS flags accordingly. The ARM flags N, Z, C, and V are defined as follows after a FMOVX FPSCR operation:

	N	Less than
15	Z	Equal
	C	Greater Than or Equal or Unordered
	V	Unordered

The format of the CDP instruction is shown in Illustration 8.

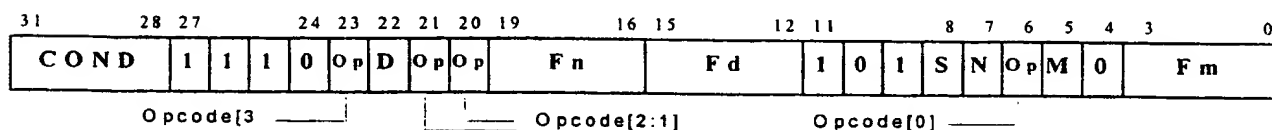


Illustration 8. CDP Instruction Format

20

Table 13. CDP Bit Field Definitions

Bit Field	Definition
Opcode	4-bit FPS operation code (See Table 14)

Bit Field	Definition
D	<p>Single precision operations:</p> <p>Destination register lsb</p> <p>Double precision operations:</p> <p>Must be set to 0</p>
Fn	<p>Single precision operations:</p> <p>Source A register upper 4 bits OR</p> <p>Extend opcode most significant 4 bits</p> <p>Double precision operations:</p> <p>Source A register address OR</p> <p>Extend opcode most significant 4 bits</p>
Fd	<p>Single precision operations:</p> <p>Destination register upper 4 bits</p> <p>Double precision operations:</p> <p>Destination register address</p>
S	<p>Operation operand size.</p> <p>0 - Single precision operands</p> <p>1 - Double precision operands</p>
N	<p>Single precision operations:</p> <p>Source A register lsb</p> <p>Extend opcode lsb</p> <p>Double precision operations:</p> <p>Must be set to 0</p> <p>Extend opcode lsb</p>
M	<p>Single precision operations:</p> <p>Source B register lsb</p> <p>Double precision operations:</p> <p>Must be set to 0</p>

<u>Bit Field</u>	<u>Definition</u>
Fm	Single precision operations: Source B register address upper 4 bits Double precision operations: Source B register address

4.8.1 Opcodes

Table 14 lists the primary opcodes for the CDP instructions. All mnemonics have the form [OPERATION][COND][S/D].

5 Table 14. CDP Opcode Specification

<u>Opcode</u>	<u>Operation</u>	
<u>Field</u>	<u>Name</u>	<u>Operation</u>
0000	FMAC	$Fd = Fn * Fm + Fd$
0001	FNMAC	$Fd = -(Fn * Fm + Fd)$
0010	FMSC	$Fd = Fn * Fm - Fd$
0011	FNMSC	$Fd = -(Fn * Fm - Fd)$
0100	FMUL	$Fd = Fn * Fm$
0101	FN MUL	$Fd = -(Fn * Fm)$
0110	FSUB	$Fd = Fn - Fm$
0111	FNSUB	$Fd = -(Fn - Fm)$
1000	FADD	$Fd = Fn + Fm$
1001 - 1011	Reserved	
1100	FDIV	$Fd = Fn / Fm$
1101	FRDIV	$Fd = Fm / Fn$
1110	FRMD	$Fd = Fn \% Fm$ (Fd = fraction left after Fn / Fm)

Opcode Field	Operation Name	Operation
1111	Extend	Use Fn register field to specify operation for 2 operand operations (See Table 15)

4.8.2 Extended Operations

Table 15 lists the extended operations available using the Extend value in the opcode field. All instructions have the form [OPERATION][COND][S/D] with the exception of the serializing and FLSCB instructions. The instruction encoding for the Extended operations is formed in the same way as the index into the register file for the Fn operand, i.e., {Fn[3:0], N}.

Table 15. CDP Extended Operations

Fn N	Name	Operation
00000	FCPY	$Fd = Fm$
00001	FABS	$Fd = \text{abs}(Fm)$
00010	FNEG	$Fd = -(Fm)$
00011	FSQRT	$Fd = \text{sqrt}(Fm)$
00100 - 00111	Reserved	
01000	FCMP*	Flags := $Fd \Leftrightarrow Fm$
01001	FCMPE*	Flags := $Fd \Leftrightarrow Fm$ with exception reporting
01010	FCMP0*	Flags := $Fd \Leftrightarrow 0$
01011	FCMPE0*	Flags := $Fd \Leftrightarrow 0$ with exception reporting

<u>Fn N</u>	<u>Name</u>	<u>Operation</u>
01100 — 01110	Reserved	
01111	FCVTD<cond>S*	Fd(double reg encoding) = Fm(single reg encoding) converted single to double precision. (coprocessor 10)
01111	FCVTS<cond>D*	Fd(single reg encoding) = Fm(double reg encoding) converted double to single precision. (coprocessor 11)
10000	FUITO*	Fd = ConvertUnsignedIntToSingle/Double(Fm)
10001	FSITO*	Fd = ConvertSignedIntToSingle/Double(Fm)
10010 - 10111	Reserved	
11000	FFTOUI*	Fd = ConvertToUnsignedInteger(Fm) {Current RMODE}
11001	FFTOUIZ*	Fd = ConvertToUnsignedInteger(Fm) {RZ mode}
11010	FFTOSI*	Fd = ConvertToSignedInteger(Fm) {Current RMODE}
11011	FFTOSIZ*	Fd = ConvertToSignedInteger(Fm) {RZ mode}
11100 - 11111	Reserved	

* Non-vectorizable operations. The LEN field is ignored and a scalar operation is

performed on the specified registers.

5. System Registers

5.1 System ID Register (FPSID)

- 5 The FPSID contains the FPS architecture and implementation-defined identification

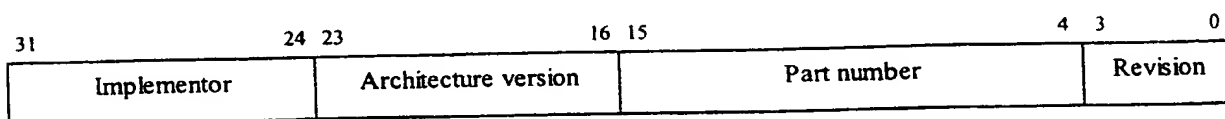


Illustration 9. FPSID Register Encoding

value. This word may be used to determine the model, feature set and revision of the FPS and the mask set number. The FPSID is read only and writes to the FPSID are ignored. See Illustration 9 for the FPSID register layout.

5.2 User Status and Control Register (FPSCR)

- 10 The FPSCR register contains user accessible configuration bits and the exception status bits. The configuration options include the exception enable bits, rounding control, vector stride and length, handling of denormal operands and results, and the use of debug mode. This register is for user and operating system code to configure

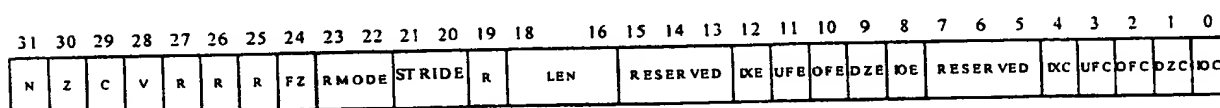


Illustration 10. User Status and Control Register (FPSCR)

- 15 the FPS and interrogate the status of completed operations. It must be saved and restored during a context switch. Bits 31 through 28 contain the flag values from the most recent compare instruction, and may be accessed using a read of the FPSCR. The FPSCR is shown in Illustration 10.

5.2.1 Compare Status and Processing Control Byte

- 20 Bits 31 through 28 contain the result of the most recent compare operation and several control bits useful in specifying the arithmetic response of the FPS in special circumstances. The format of the Compare Status and Processing Control Byte are

given in Illustration 11.

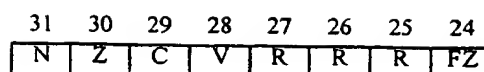


Illustration 11. FPSCR Compare Status and Processing
Control Byte

Table 16. FPSCR Compare Status and Processing Control Byte Field Definitions

Register Bit	Name	Function
31	N	Compare result was less than
30	Z	Compare result was equal
29	C	Compare result was greater than or equal or unordered
28	V	Compare result was unordered
27:25	Reserved	
24	FZ	Flush to zero 0 : IEEE 754 Underflow handling (Default) 1 : Flush tiny results to zero Any result which is smaller than the normal range for the destination precision will result in a zero written to the destination. The UNDERFLOW exception trap will not be taken.

5 5.2.2 System Control Byte

The system control byte controls the rounding mode, vector stride and vector length fields. The bits are specified as shown in Illustration 12.

The VFPv1 architecture incorporates a register file striding mechanism for use with vector operations. If the STRIDE bits are set to 00, the next register selected in a vector operation will be the register immediately following the previous register in the register file. The normal register file wrapping mechanism is unaffected by the stride value. A STRIDE of 11 will increment all input registers and the output register by 2.

For example,

FMULEQS F8, F16, F24

10

will perform the following non-vector operations:

FMULEQS F8, F16, F24

FMULEQS F10, F18, F26

FMULEQS F12, F20, F28

15 FMULEQS F14, F22, F30

effectively 'striding' the operands for the multiply in the register file by 2 rather than by 1 register

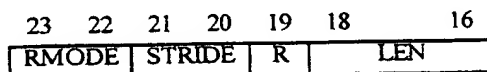


Illustration 12. FPSCR System

Control Byte

20

Table 17. FPSCR System Control Byte Field Definitions

Register Bit	Name	Function
23:22	RMODE	Set rounding mode 00 : RN (Round to Nearest, Default) 01 : RP (Round towards Plus Infinity) 10 : RM (Round towards Minus Infinity) 11 : RZ (Round towards Zero)
21:20	STRIDE	Set the vector register access to: 00 : 1 (Default) 01 : RESERVED 10 : RESERVED 11 : 2
19	Reserved (R)	
18:16	LEN	Vector Length. Specifies length for vector operations. (Not all encodings are available in each implementation.) 000 : 1 (Default) 001 : 2 010 : 3 011 : 4 100 : 5 101 : 6 110 : 7 111 : 8

5.2.3 Exception Enable Byte

The exception enable byte occupies bits 15:8 and contains the enables for exception traps. The bits are specified as shown in Illustration 13. The exception enable bits conform to the requirements of the IEEE 754 specification for handling of floating point exception conditions. If the bit is set, the exception is enabled, and FPS will

5 signal a user visible trap to the operating system in the event of an occurrence of the exceptional condition on the current instruction. If the bit is cleared, the exception is not enabled, and the FPS will not signal a user visible trap to the operating system in the event of the exceptional condition, but will generate a mathematically reasonable result. The default for the exception enable bits is disabled. For more information on

10 exception handling please see the IEEE 754 standard.

Some implementations will generate a bounce to the support code to handle exceptional conditions outside the capability of the hardware, even when the exception is disabled. This will be generally invisible to user code.

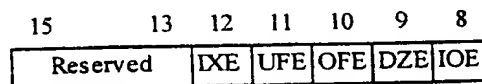


Illustration 13. FPSCR Exception Enable

15 Table 18. FPSCR Exception Enable Byte Fields

Register Bit	Name	Function
15:13	Reserved	
12	IXE	Inexact Enable Bit 0: Disabled (Default) 1: Enabled
11	UFE	Underflow Enable Bit 0: Disabled (Default) 1: Enabled

Register Bit	Name	Function
10	OFE	Overflow Enable Bit 0: Disabled (Default) 1: Enabled
9	DZE	Divide-by-Zero Enable Bit 0: Disabled (Default) 1: Enabled
8	IOE	Invalid Operand Enable Bit 0: Disabled (Default) 1: Enabled

5.2.4 Exception Status Byte

The exception status byte occupies bits 7:0 of the FPSCR and contains the exception status flag bits. There are five exception status flag bits, one for each floating point exception. These bits are 'sticky'; once set by a detected exception, they must be cleared by a FMOVX write to the FPSCR or a FSERIALCL instruction. The bits are specified as shown in Illustration 14. In the case of an enabled exception, the corresponding exception status bit will not be automatically set. It is the task of the support code to set the proper exception status bit as needed. Some exceptions may be automatic, i.e., if the exception condition is detected, the FPS will bounce on the subsequent floating point instruction regardless of how the exception enable bit is set. This allows some of the more involved exception processing required by the IEEE 754 standard to be performed in software rather than in hardware. An example would be an underflow condition with the FZ bit set to 0. In this case, the correct result may be a denormalized number depending on the exponent of the result and the rounding mode. The FPS allows implementers to select the response including the option to bounce and utilize the support code to produce the correct result and write this value to the destination register. If the underflow exception enable bit is set, the user's trap

handler will be called after the support code has completed the operation. This code may alter the state of the FPS and return, or terminate the process.

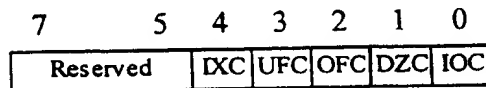


Illustration 14. FPSCR Exception

Status Byte

5 Table 19. FPSCR Exception Status Byte Field Definitions

Register Bit	Name	Function
7:5	Reserved	
4	IXC	Inexact exception detected
3	UFC	Underflow exception detected
2	OFC	Overflow exception detected
1	DZC	Divide by zero exception detected
0	IOC	Invalid Operation exception detected

5.3 Register File Content Register (FPREG)

The Register File Content Register is a privileged register containing information which may be used by a debugger to properly present the contents of the register as interpreted by the currently running program. The FPREG contains 16 bits, one bit

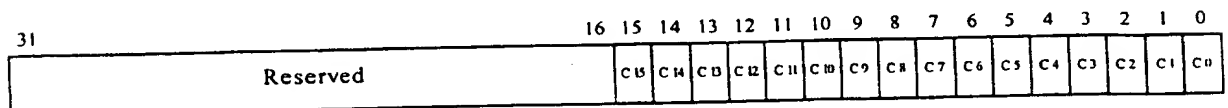


Illustration 15. FPREG Register Encoding

for each double precision register in the register file. If the bit is set, the physical register pair represented by the bit is to be displayed as a double precision register. If the bit is clear, the physical register is uninitialized or contains one or two single

precision data values.

Table 20. FPREG Bit Field Definitions

FPREG bit	Bit Set	Bit Clear
C0	D0 valid	S1 and S0 valid or uninitialized
C1	D1 valid	S3 and S2 valid or uninitialized
C2	D2 valid	S5 and S4 valid or uninitialized
C3	D3 valid	S7 and S6 valid or uninitialized
C4	D4 valid	S9 and S8 valid or uninitialized
C5	D5 valid	S11 and S10 valid or uninitialized
C6	D6 valid	S13 and S12 valid or uninitialized
C7	D7 valid	S15 and S14 valid or uninitialized
C8	D8 valid	S17 and S16 valid or uninitialized
C9	D9 valid	S19 and S18 valid or uninitialized
C10	D10 valid	S21 and S20 valid or uninitialized
C11	D11 valid	S23 and S22 valid or uninitialized

FPREG bit	Bit Set	Bit Clear
C12	D12 valid	S25 and S24 valid or uninitialized
C13	D13 valid	S27 and S26 valid or uninitialized
C14	D14 valid	S29 and S28 valid or uninitialized
C15	D15 valid	S31 and S30 valid or uninitialized

6. Exception Processing

- 5 The FPS operates in one of two modes, a debug mode and a normal mode. If the DM bit is set in the FPSCR, the FPS operates in debug mode. In this mode the FPS executes one instruction at a time while ARM is made to wait until the exception status of the instruction is known. This will guarantee the register file and memory are precise with respect to instruction flow, but at the expense of much increased
- 10 execution time. The FPS will accept a new instruction from the ARM when resources allow, and signal exceptions upon detection of the exceptional condition. Exception reporting to the ARM will always be precise with respect to the floating point instruction stream except in the case of a load or store operation which follows a vector operation and executes in parallel with the vector operation. In this case the
- 15 contents of the register file, for load operations, or memory, for store operations, may not be precise.

6.1 Support Code

- Implementations of the FPS may elect to be IEEE 754 compliant with a combination of hardware and software support. For unsupported data types and automatic
- 20 exceptions, the support code will perform the function of compliant hardware and

return the result, when appropriate, to the destination register and return to the user's code without calling a user's trap handler or otherwise modifying the flow of the user's code. It will appear to the user that the hardware alone was responsible for the processing of the floating point code. Bouncing to support code to handle these
5 features significantly increases the time to perform or process the feature, but the incidence of these situations is typically minimal in user code, embedded applications, and well written numeric applications.

The support code is intended to have two components: a library of routines which
10 perform operations beyond the scope of the hardware, such as transcendental computations, as well as supported functions, such as divide with unsupported inputs or inputs which may generate an exception; and a set of exception handlers which process exception traps in order to provide IEEE 754 compliance. The support code is required to perform implemented functions in order to emulate proper handling of any
15 unsupported data type or data representation (e.g., denormal values). The routines may be written to utilize the FPS in their intermediate calculations if care is taken to restore the users' state at the exit of the routine.

6.2 Exception Reporting and Processing

Exceptions in normal mode will be reported to the ARM on the next floating point
20 instruction issued after the exception condition is detected. The state of the ARM processor, the FPS register file, and memory may not be precise with respect to the offending instruction at the time the exception is taken. Sufficient information is available to the support code to correctly emulate the instruction and process any exception resulting from the instruction.

25

In some implementations, support code may be used to process some or all operations with special IEEE 754 data, including infinities, NaNs, denormal data, and zeros. Implementations which do so will refer to these data as unsupported, and bounce to the support code in a manner generally invisible to user code, and return with the

IEEE 754 specified result in the destination register. Any exceptions resulting from the operation will abide by the IEEE 754 rules for exceptions. This may include trapping to user code if the corresponding exception enable bit is set.

- 5 The IEEE 754 standard defines the response to exceptional conditions for both cases of the exception enabled and disabled in the FPSCR. The VFPv1 Architecture does not specify the boundary between the hardware and software used to properly comply with the IEEE 754 specification.

10 **6.2.1 Unsupported Operations and Formats**

The FPS does not support any operations with decimal data or conversion to or from decimal data. These operations are required by the IEEE 754 standard and must be provided by the support code. Any attempt to utilize decimal data will require library routines for the desired functions. The FPS has no decimal data type and cannot be

- 15 used to trap instructions which use decimal data.

6.2.2 Use of FMOVX When the FPS is Disabled or Exceptional

The FMOVX instruction, executed in SUPERVISOR or UNDEFINED mode may read and write the FPSCR or read the FPSID or FPREG when the FPS is in an exceptional state or is disabled (if the implementation supports a disable option)

- 20 without causing an exception to be signalled to the ARM.

Although particular embodiments of the invention have been described herewith, it will be apparent that the invention is not limited thereto, and that many modifications and additions may be made within the scope of the invention. For example, various combinations of the features of the following dependent claims could be made with the

25 features of the independent claims without departing from the scope of the present invention.

CLAIM

1. An apparatus for processing data, said apparatus comprising:
 - a register bank having a plurality of registers; and
 - 5 an instruction decoder responsive to at least one data processing instruction specifying a vector operation that executes a data processing operation a plurality of times using data values from a sequence of registers within said register bank; wherein
 - said register bank includes at least one subset of registers, said sequence of
 - 10 registers being within said subset; and
 - said instruction decoder controls said sequence of registers to wrap within said subset of registers.
2. An apparatus as claimed in claim 1, wherein
 - 15 said vector operation executes said data processing operation using a plurality of respective data values from a corresponding plurality of sequences of registers;
 - said register bank contains a plurality of subsets of registers, said plurality of sequences of registers being within respective subsets; and
 - said instruction decoder controls said plurality of sequences of registers to
 - 20 wrap within respective subsets of registers.
3. Apparatus as claimed in claim 2, wherein said plurality of subsets are disjoint.
4. Apparatus as claimed in any one of claims 1, 2 and 3, wherein said subset
- 25 comprises a range of consecutively numbers registers.
5. Apparatus as claimed in claim 2, wherein each of said plurality of subsets comprises a range of consecutively numbered registers.

6. Apparatus as claimed in claim 5, where said plurality of subsets comprise respective contiguous ranges of consecutively numbered registers.
7. Apparatus as claimed in claim 6, comprising four contiguous ranges.
- 5 8. Apparatus as claimed in any one of the preceding claims, further comprising a memory and a transfer controller for controlling transfers of data values between said memory and registers within said register bank, said transfer controller being responsive to multiple transfer instructions to transfer a sequence of data values
- 10 between said memory and a sequence of registers within said register bank.
9. Apparatus as claimed in claim 6, wherein each range is addressed via an incrementer that wraps between the end points of that range.
- 15 10. Apparatus as claimed in any one of the preceding claims, wherein said sequence is a sequence of consecutive registers.
11. Apparatus as claimed in any one of the preceding claims, wherein said register bank and said instruction decoder are part of a floating point unit.
- 20 12. A method of processing data, said method comprising the steps of:
- storing data values within a plurality of registers of a register bank; and
- in response to at least one data processing instruction specifying a vector operation, executing a data processing operation a plurality of times using data values
- 25 from a sequence of registers within said register bank; wherein
- said register bank includes at least one subset of registers, said sequence of registers being within said subset; and
- during said executing, said sequence of registers wraps within said subset of registers.

13. A method as claimed in claim 12, wherein
said vector operation executes said data processing operation using a plurality
of respective data values from a corresponding plurality of sequences of registers;
5 said register bank contains a plurality of subsets of registers, said plurality of
sequences of registers being within respective subsets; and
during said execution, said plurality of sequences of registers to wrap within
respective subsets of registers.
- 10 14. A method as claimed in claim 13, wherein data values within one sequence are tap
coefficients of a filter and data values in another sequence are signal values for filtering
by said filter.
15. A method as claimed in claim 12, wherein a plurality of vectors operations are
15 executed upon data values within said plurality of sequences with a starting point of at
least one sequence changing with each vector operation.

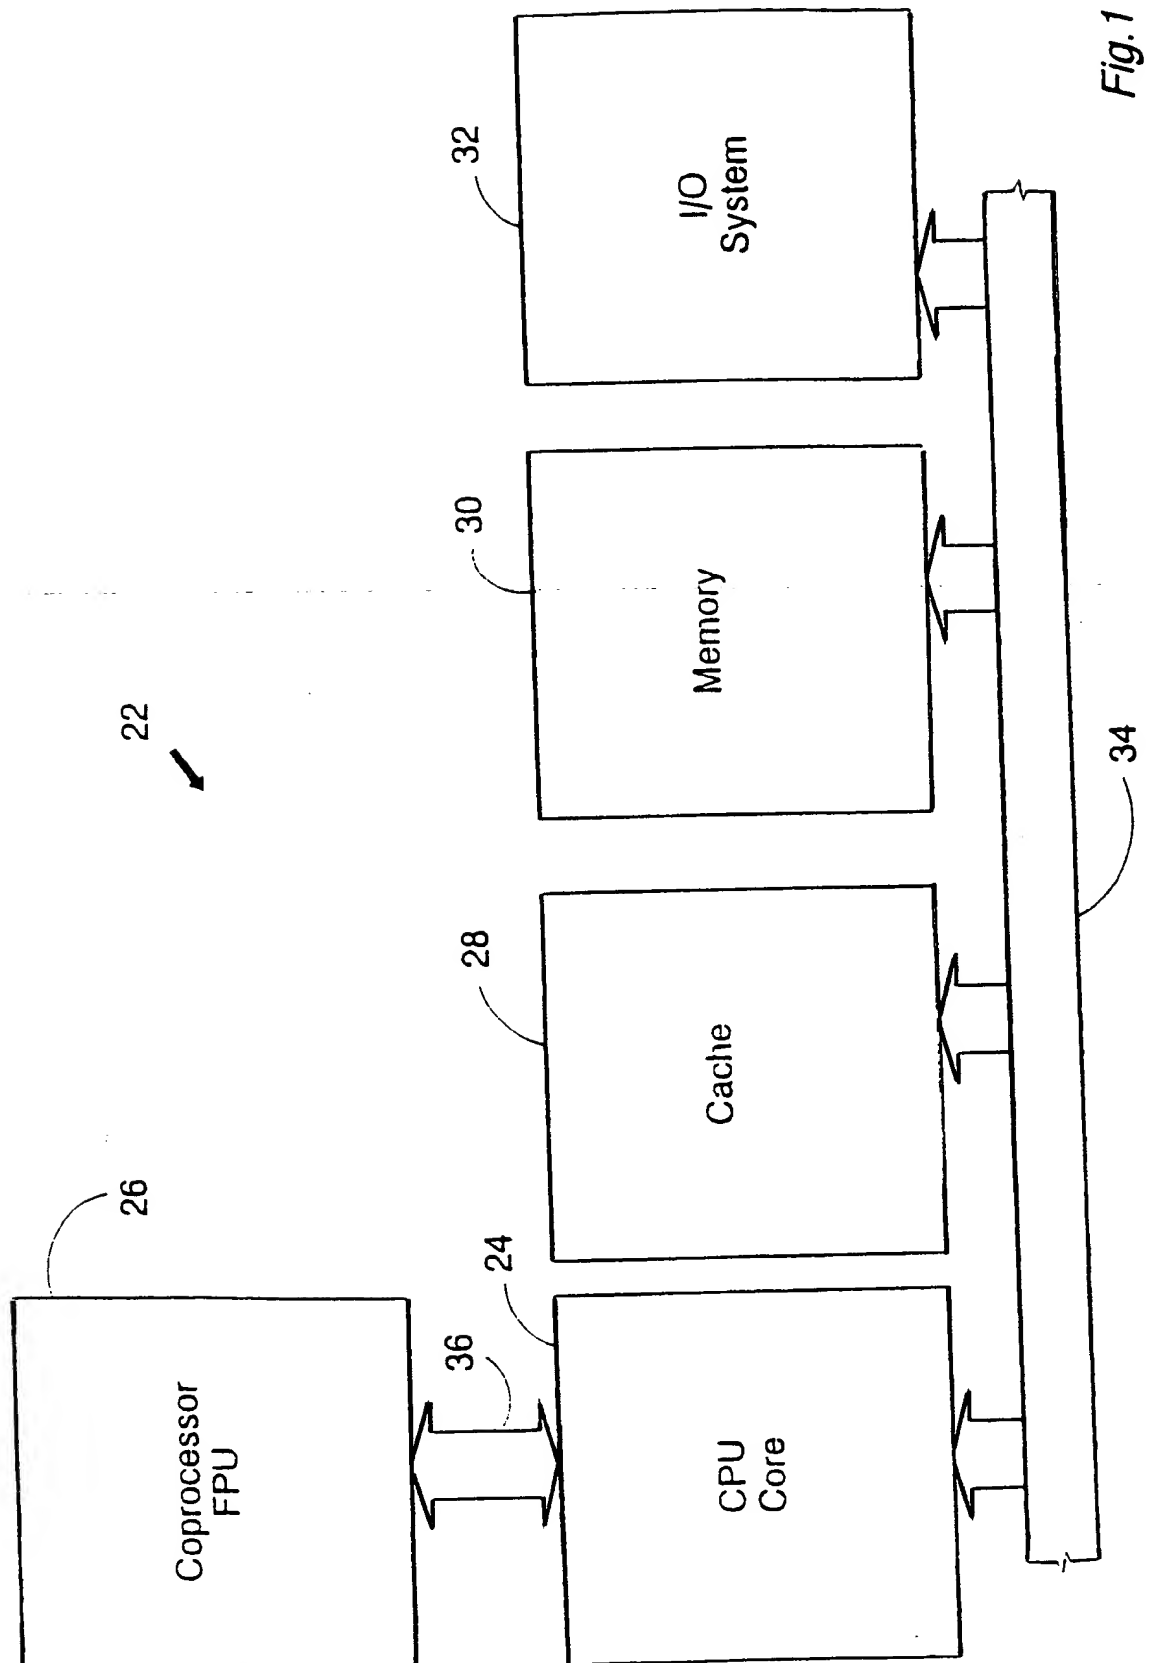


Fig.1

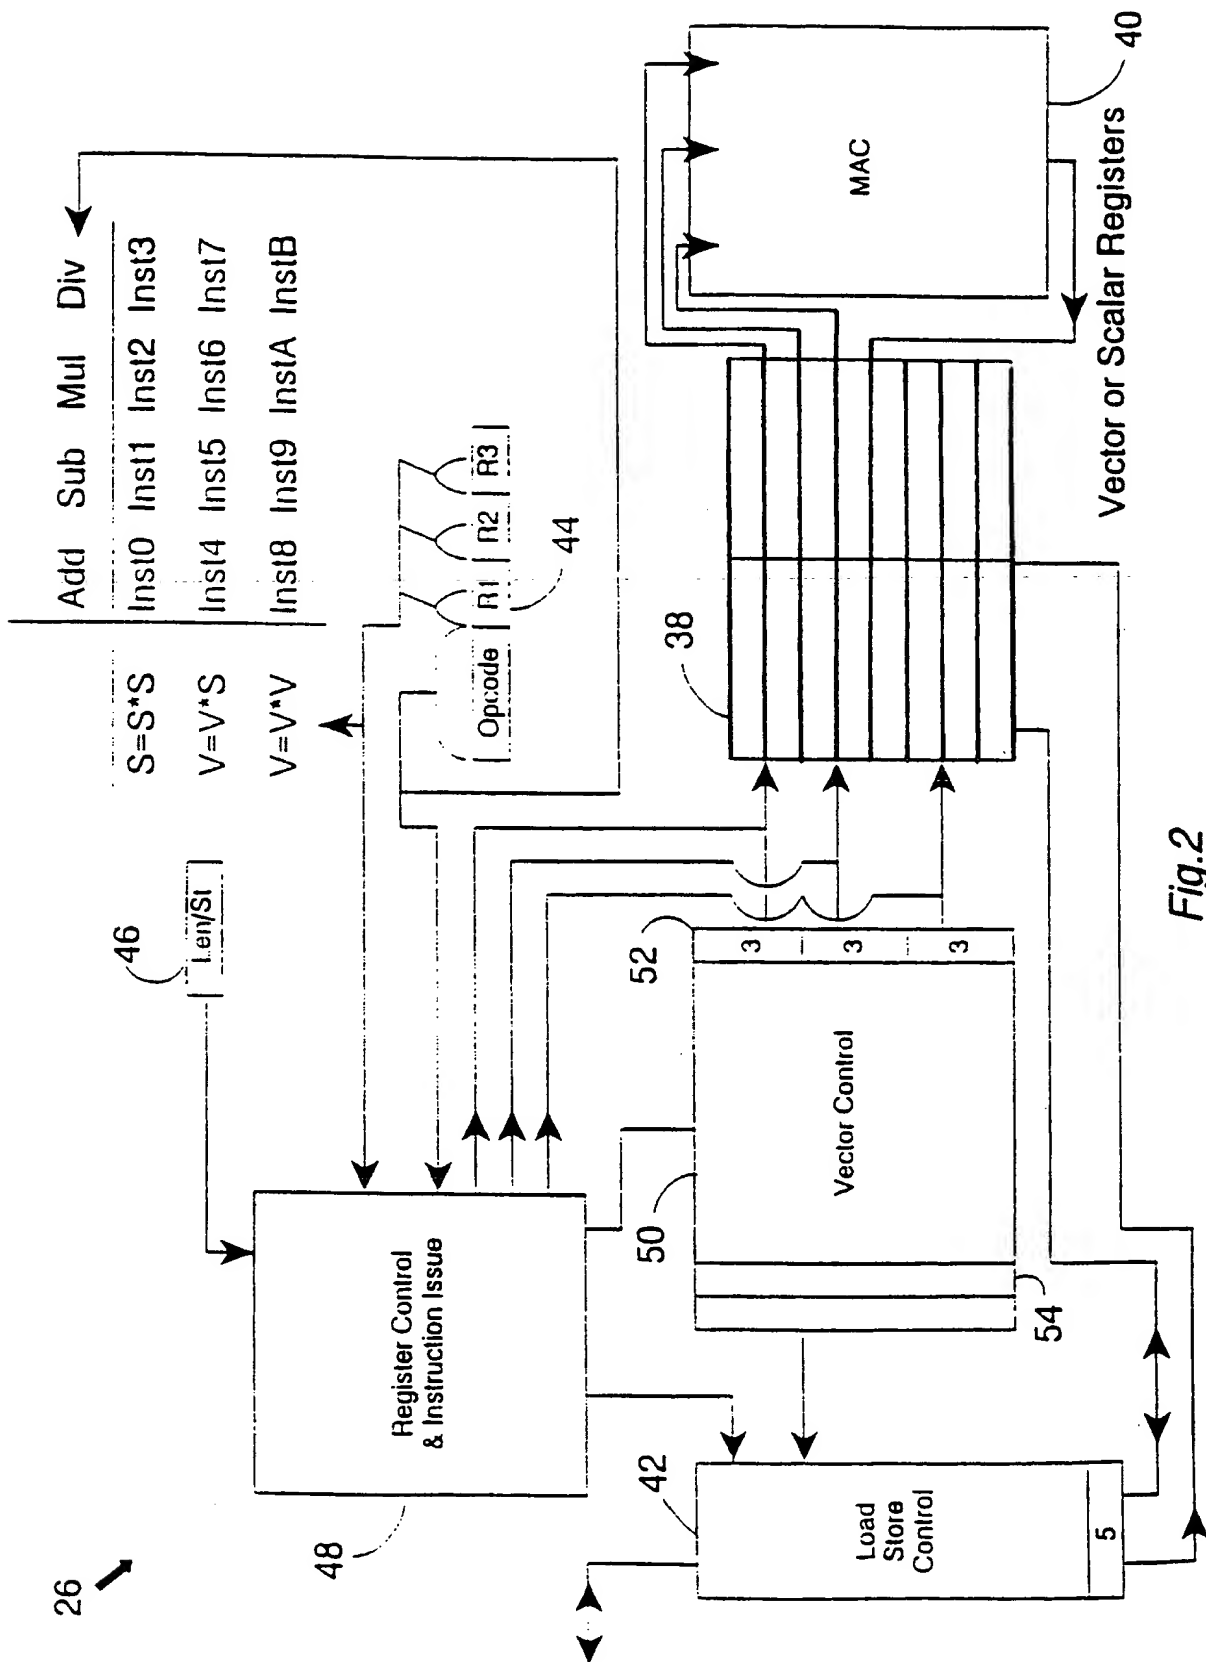
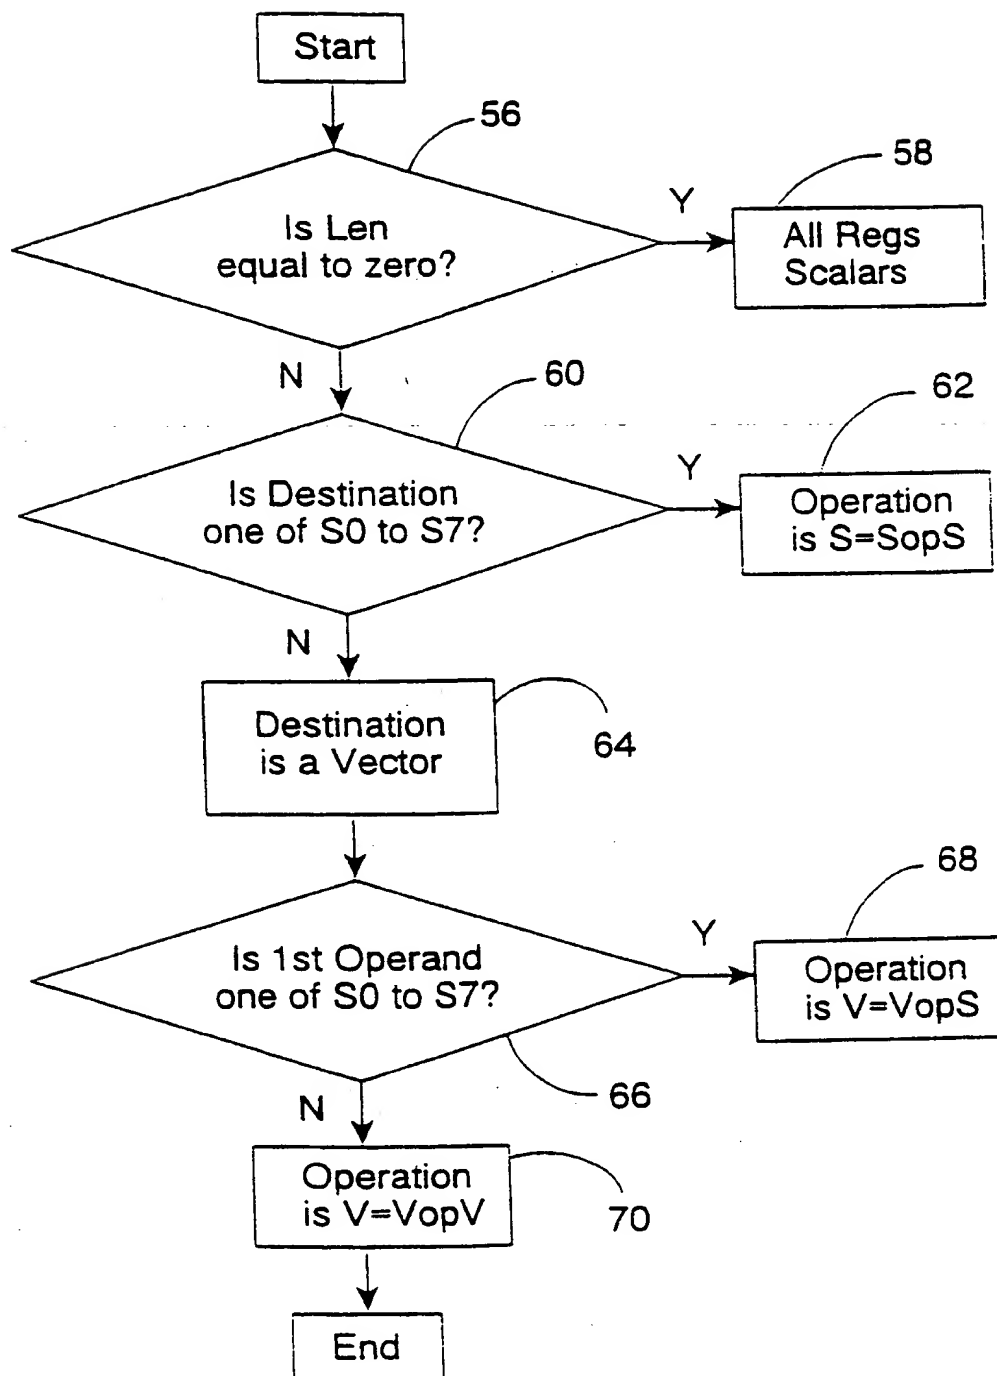
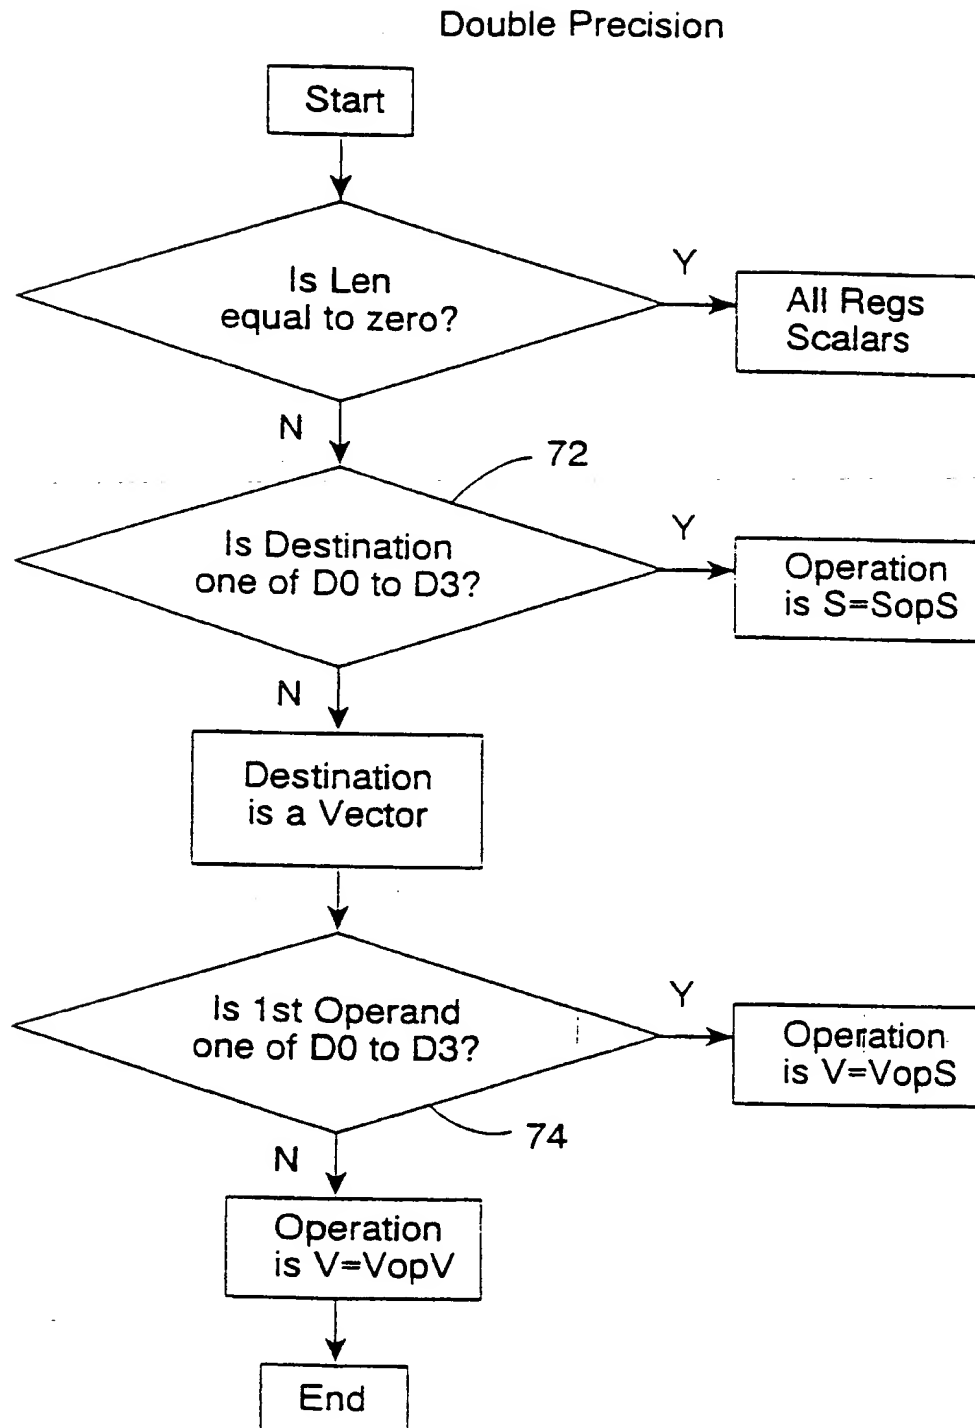


Fig.2

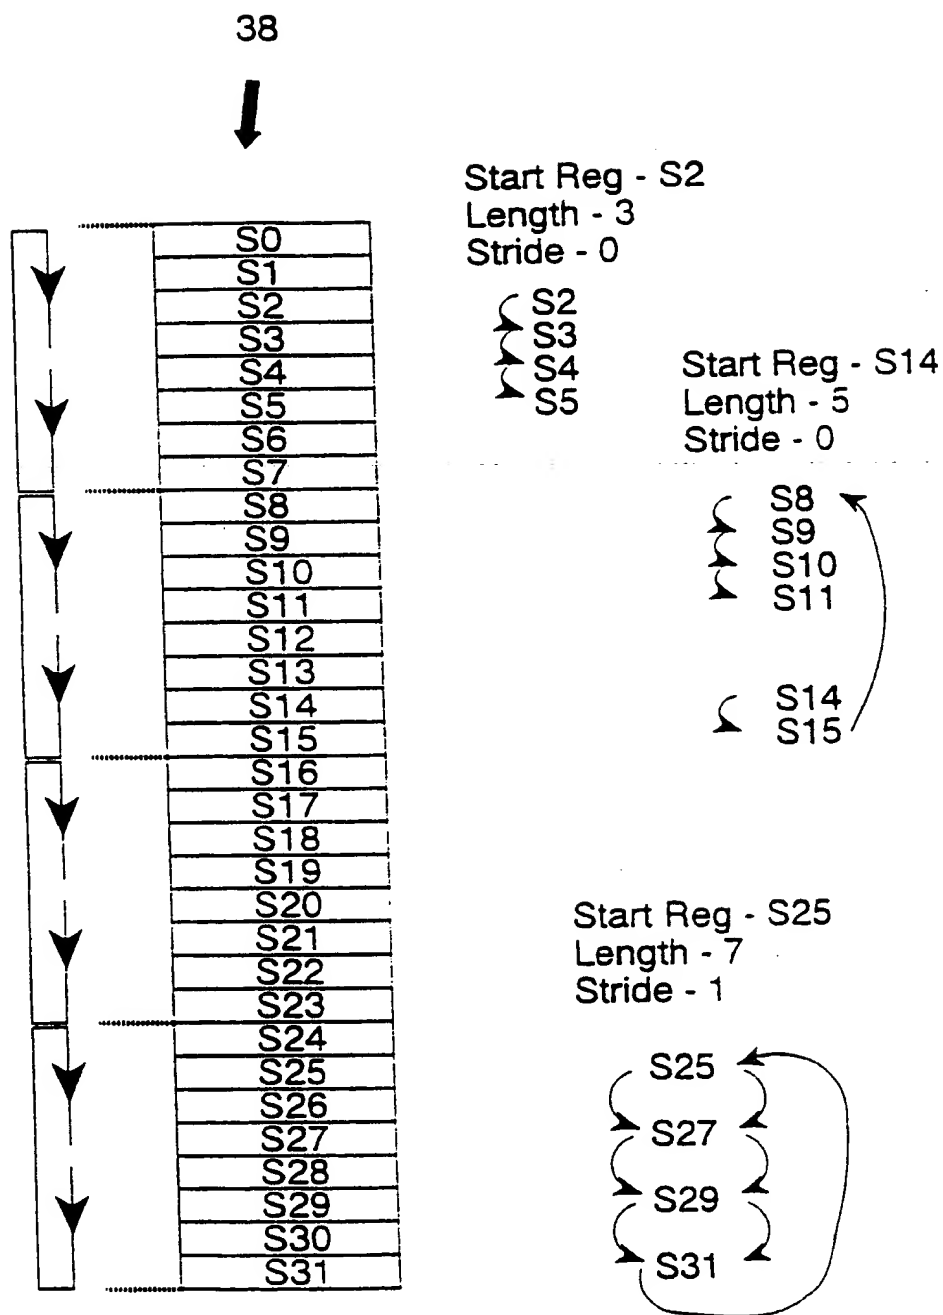
3/15

Single Precision

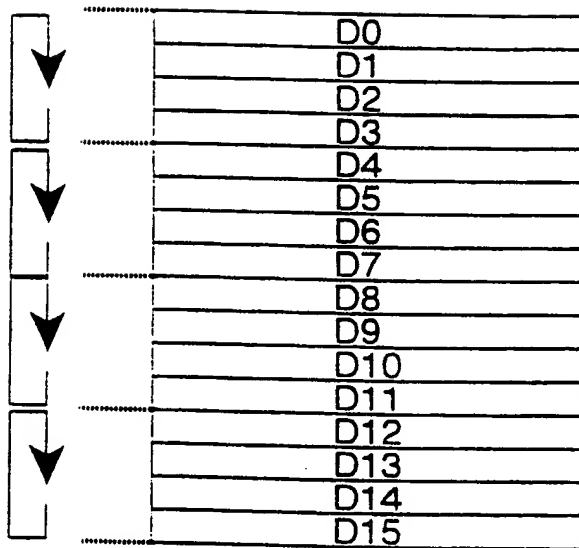




5/15



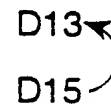
38



Start Reg - D0
Length - 3
Stride - 0



Start Reg - D15
Length - 1
Stride - 1



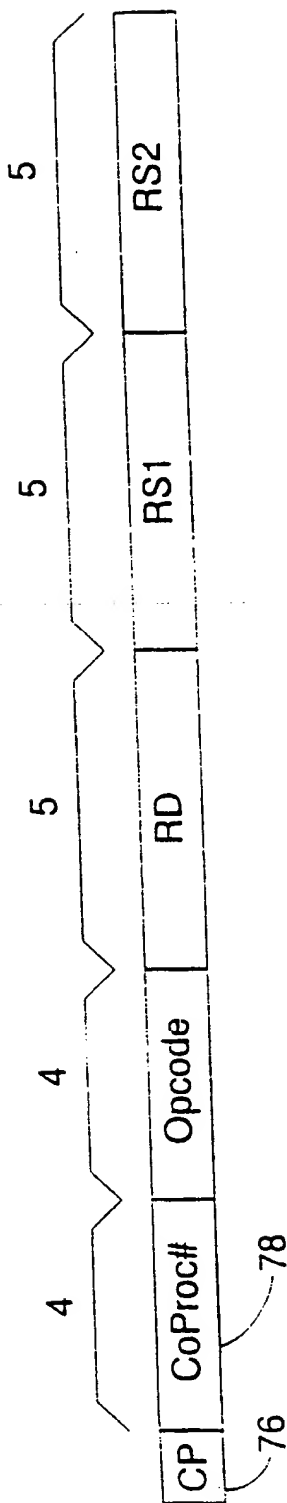


Fig. 7A

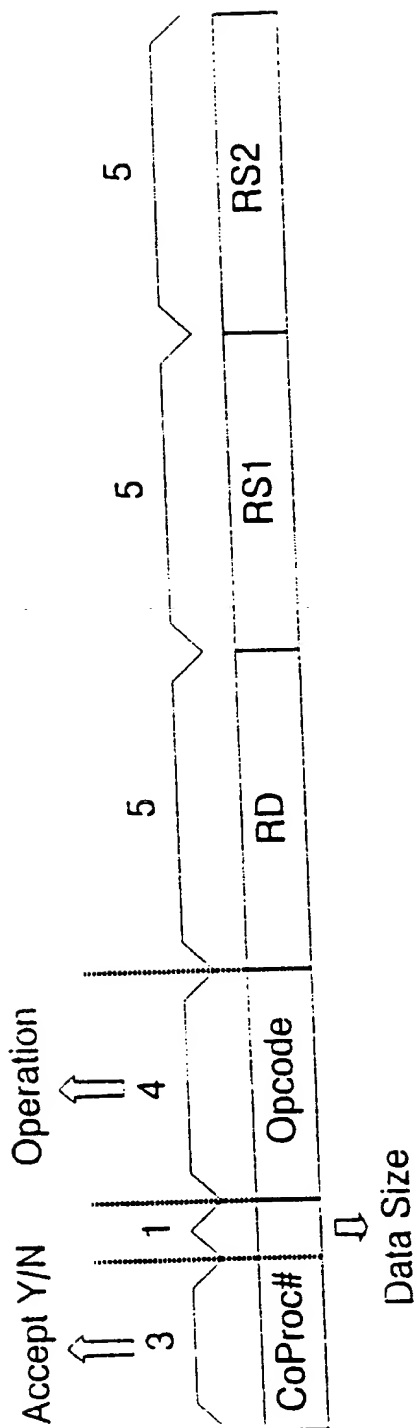


Fig. 7B

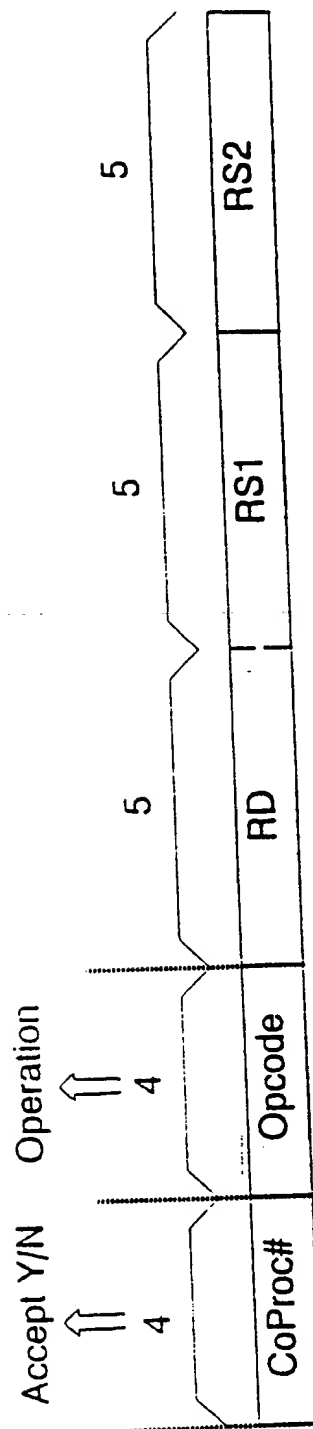


Fig. 7C

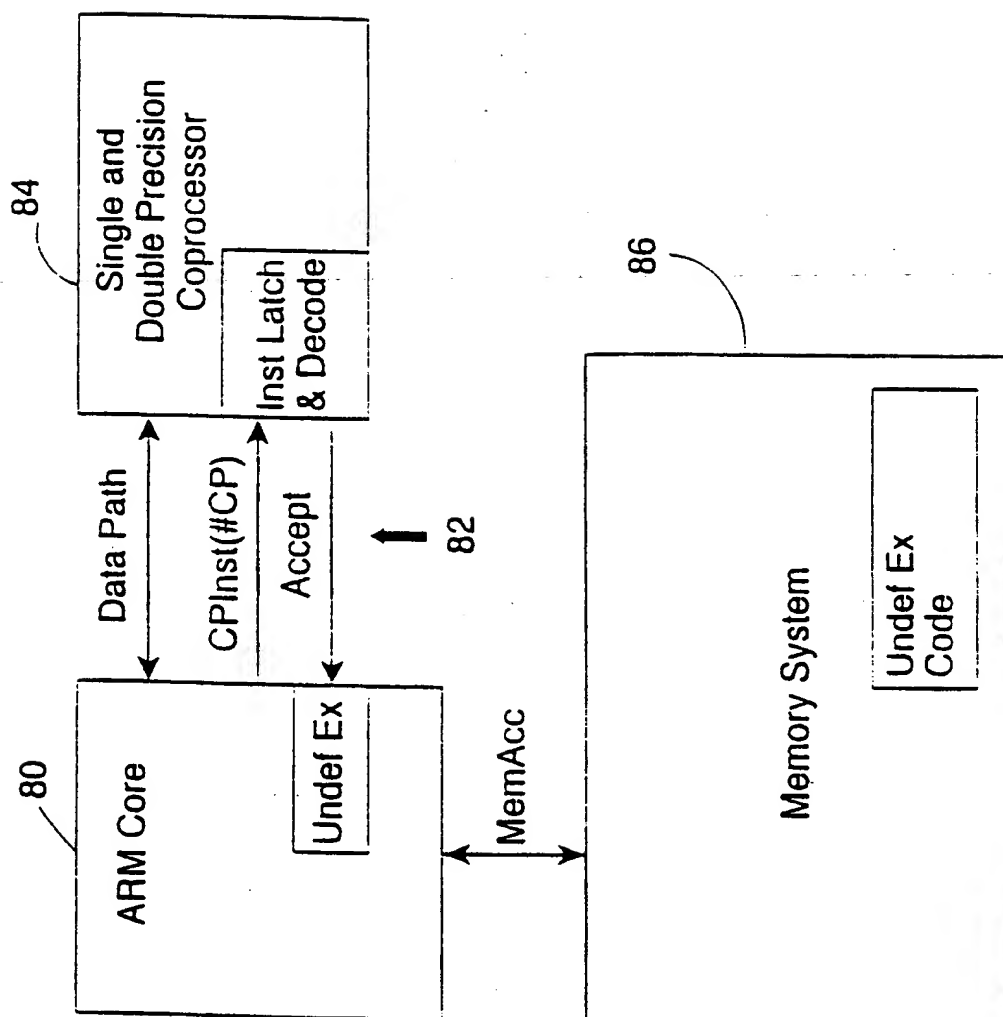


Fig.8

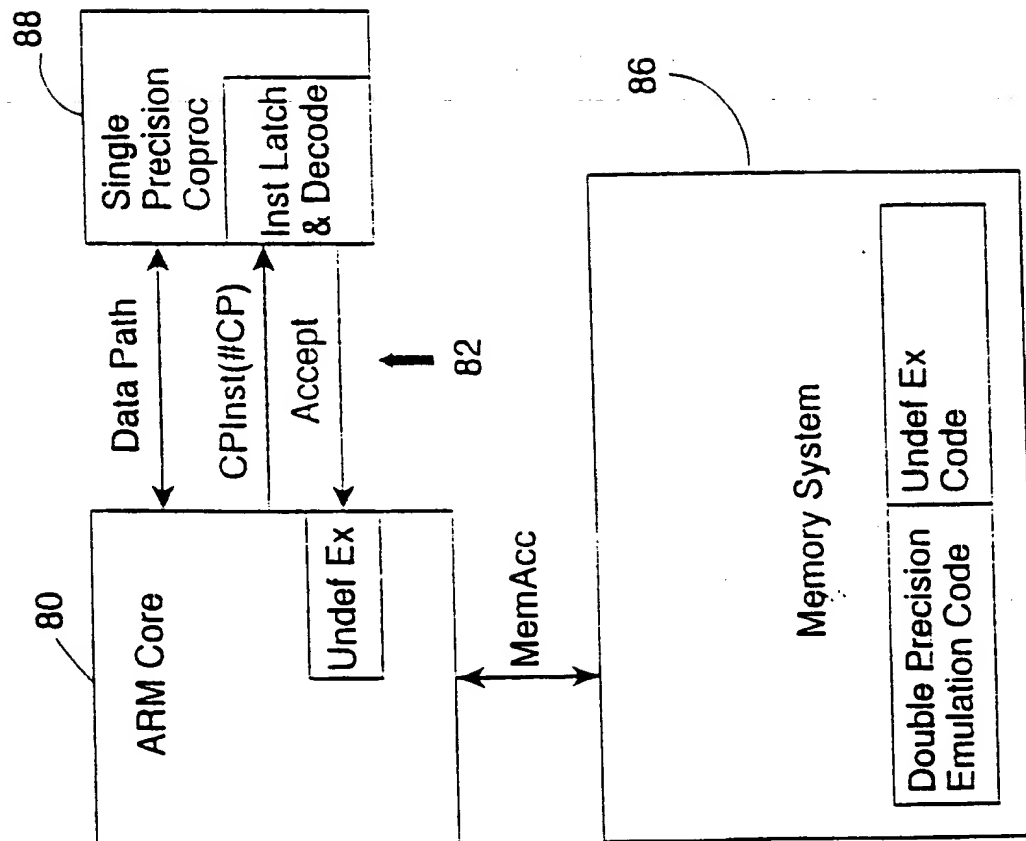


Fig.9

10/15

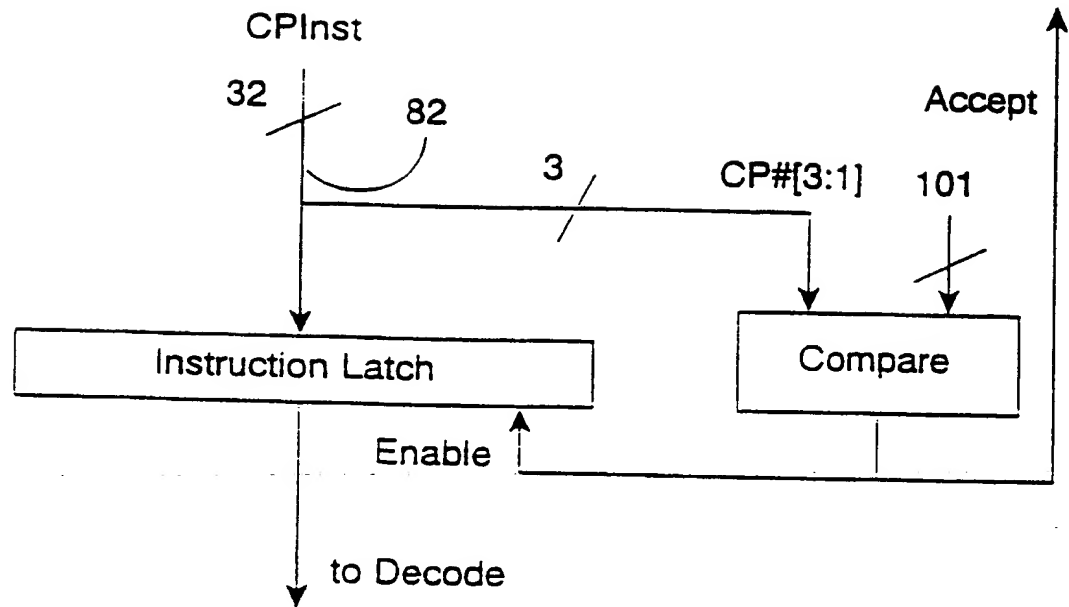


Fig. 10

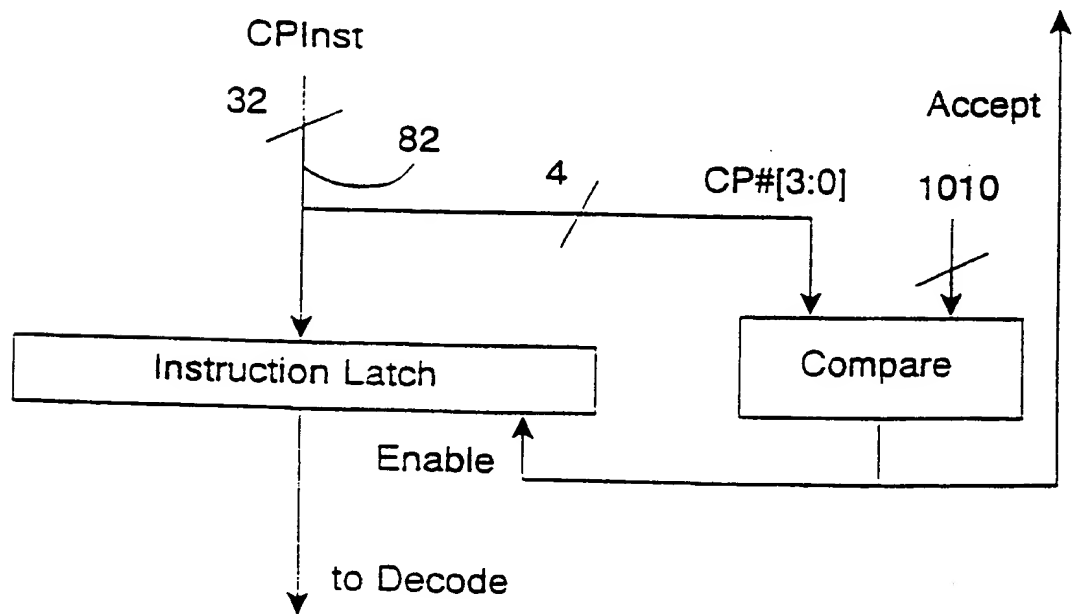


Fig. 11

SUBSTITUTE SHEET (RULE 26)

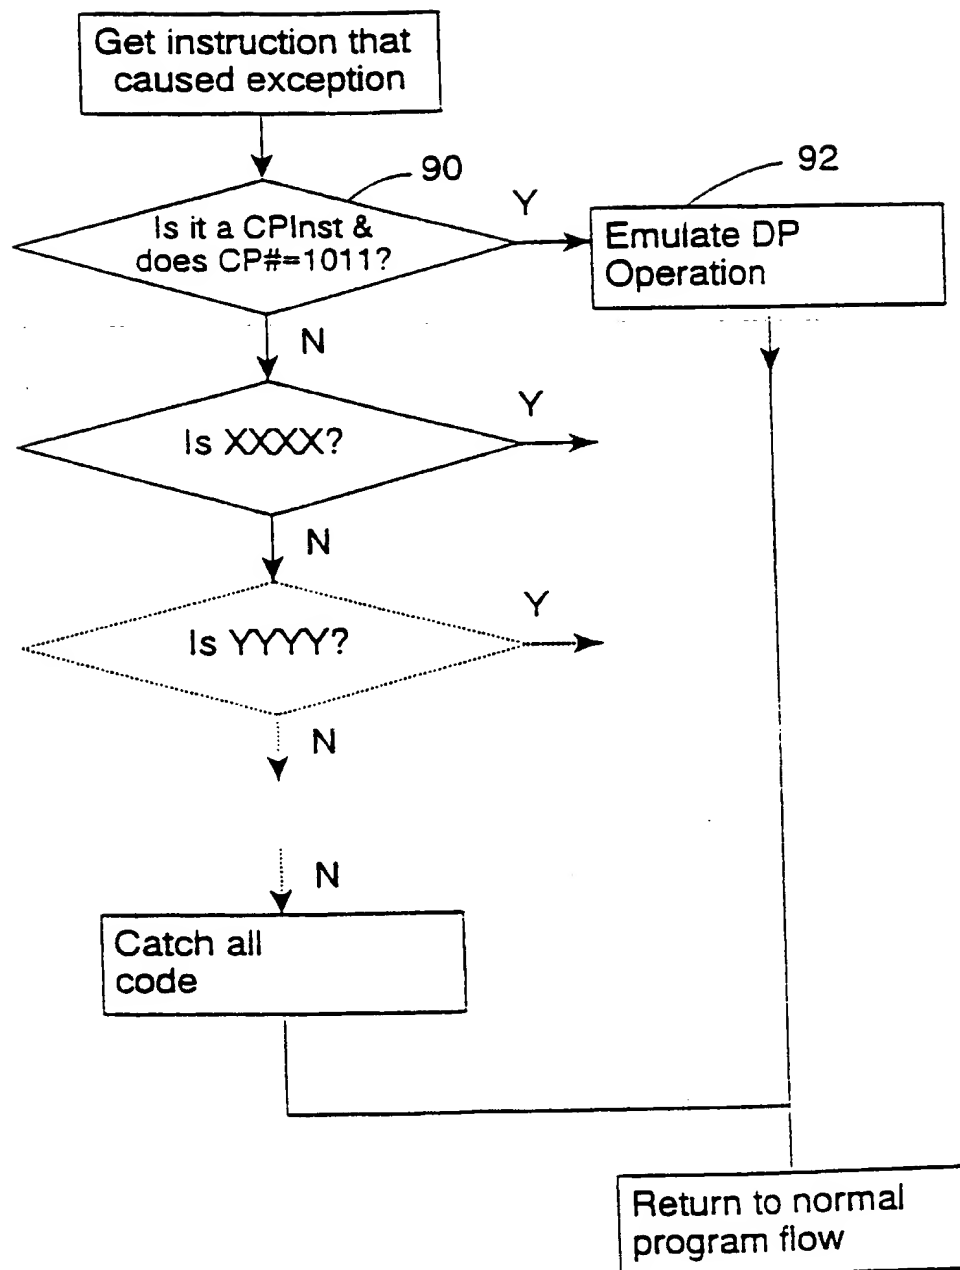


Fig. 12

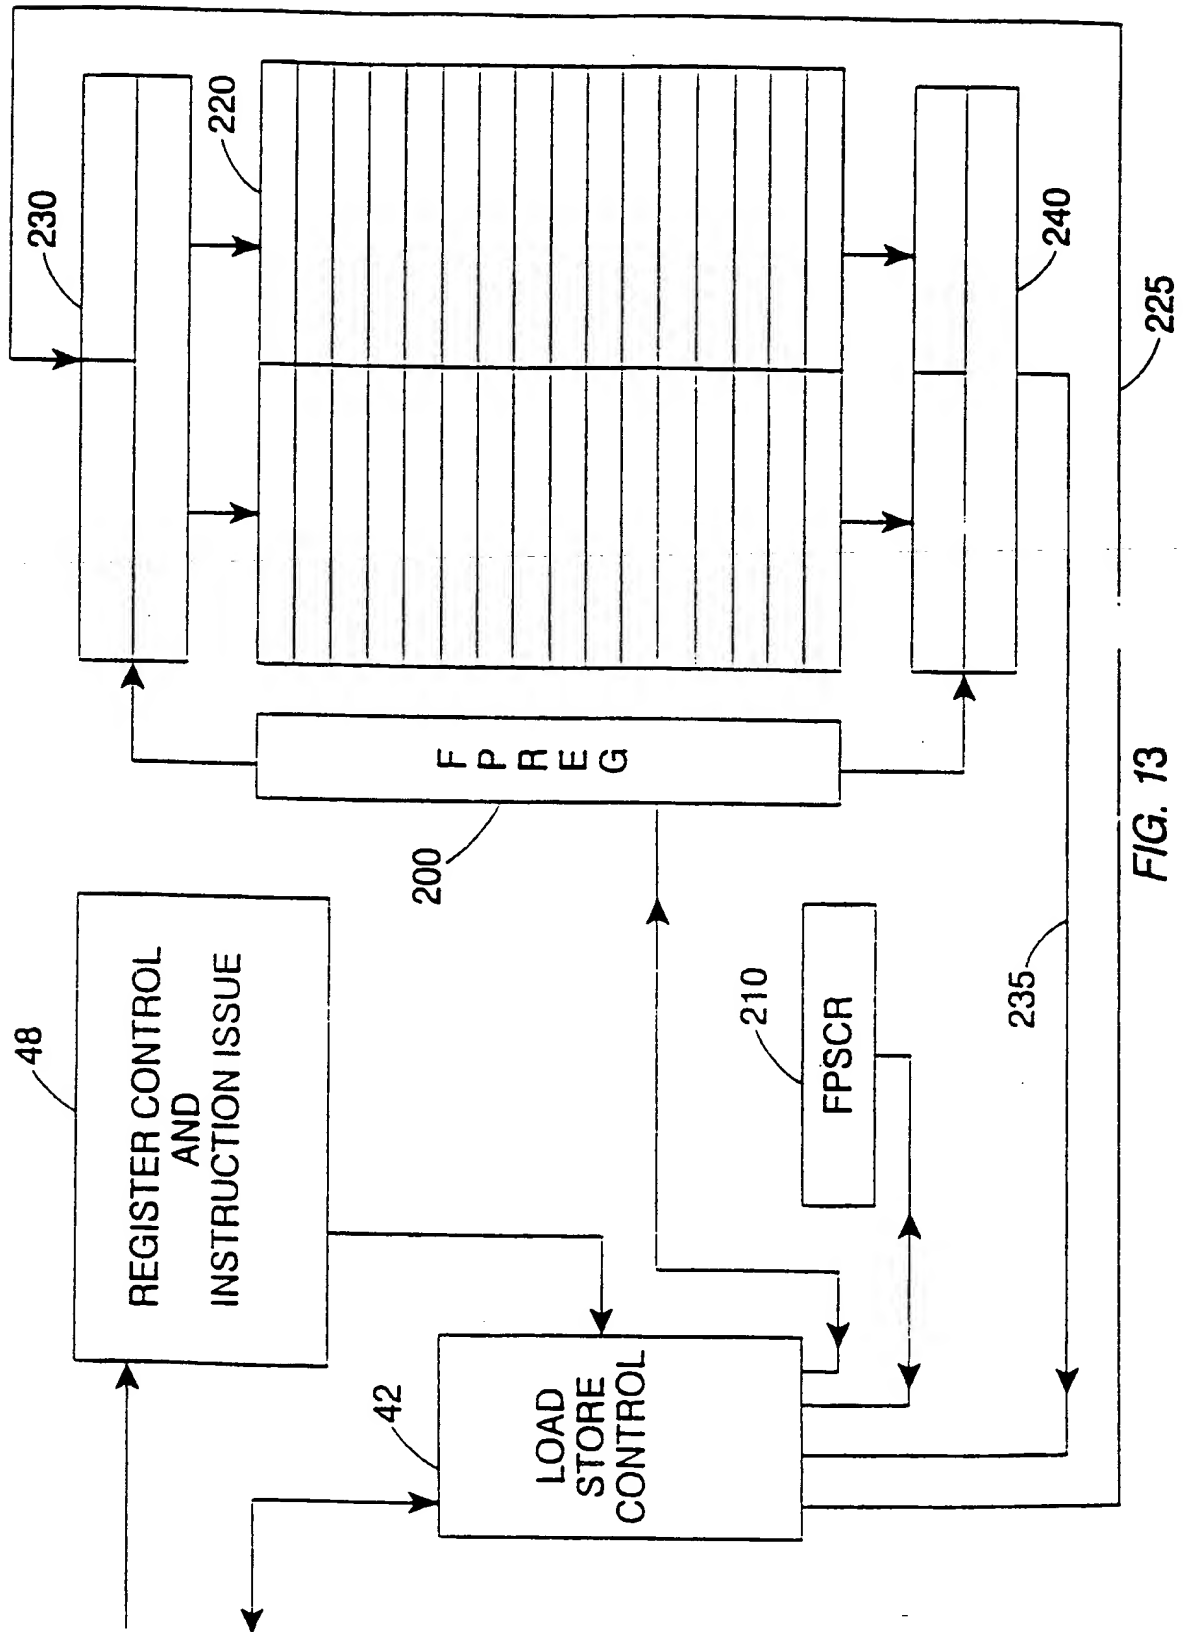


FIG. 13

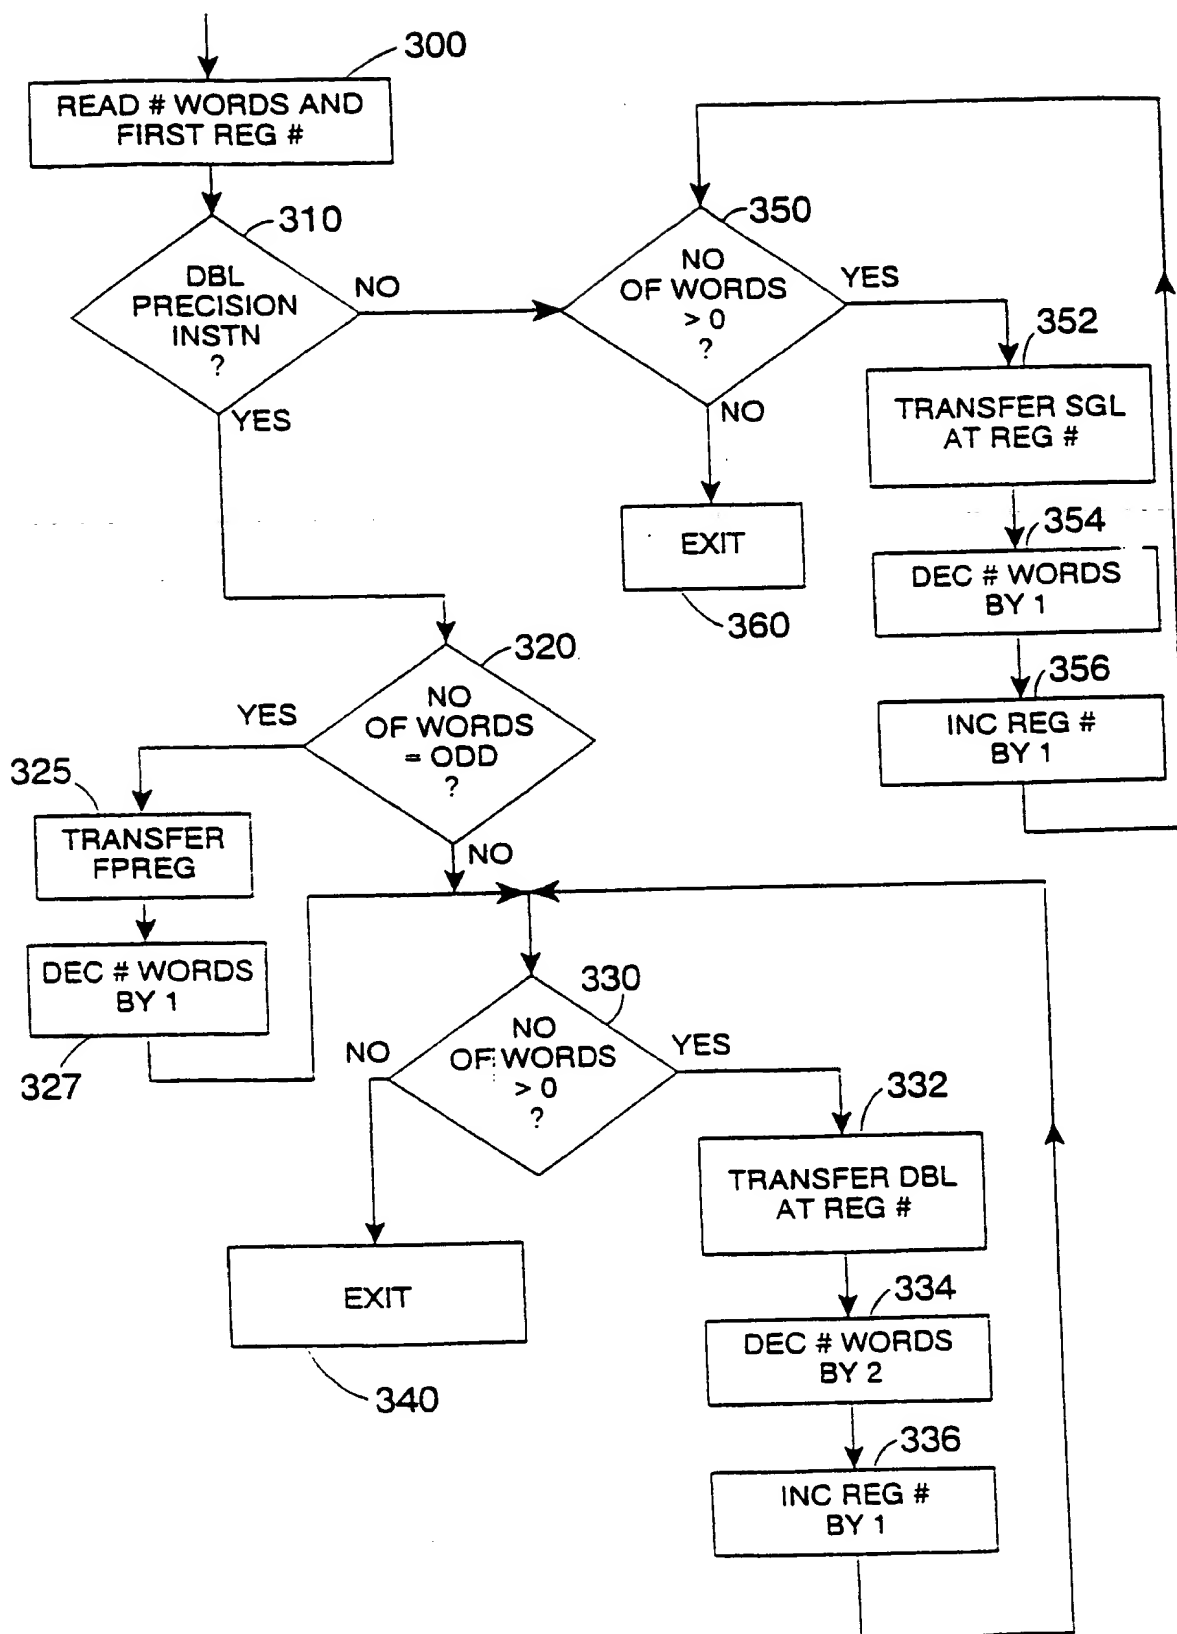
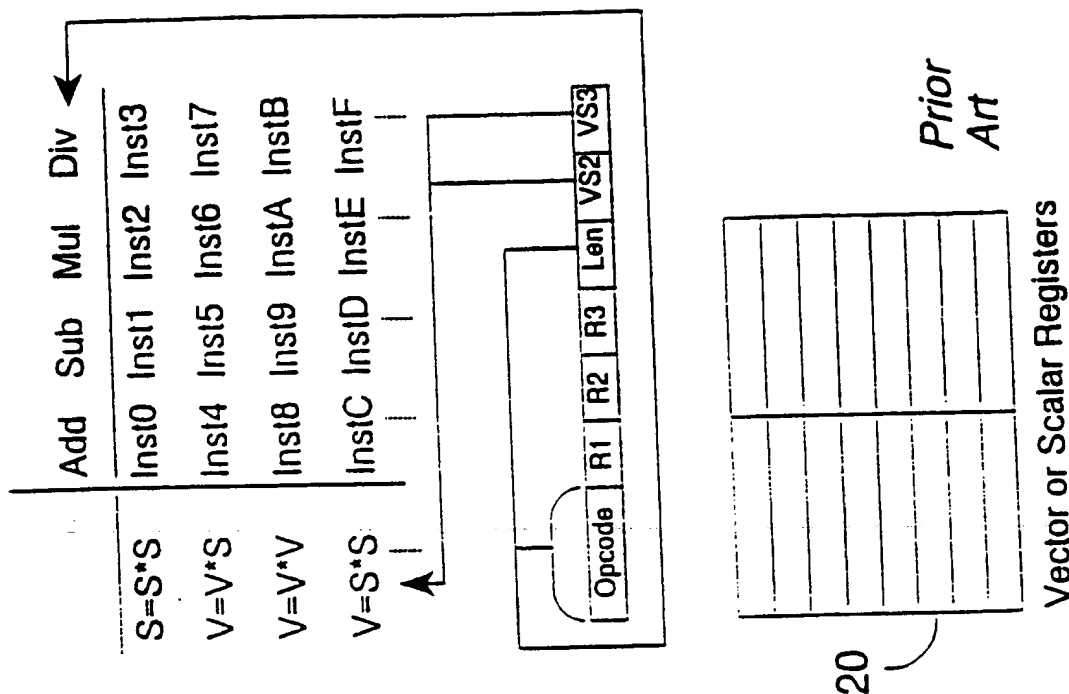


FIG. 14
SUBSTITUTE SHEET (RULE 26)

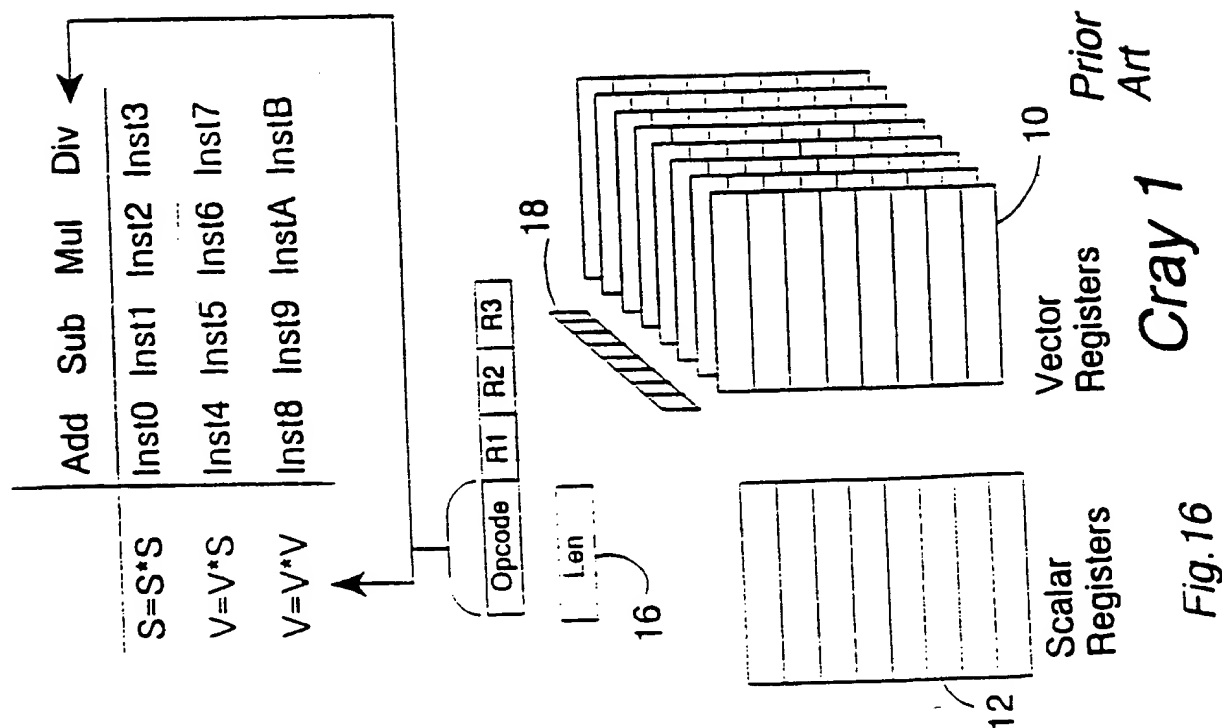
1	DH	DL
0	S	S
0	S	S
0	S	S
1	DH	DL
0	S	S

FIG. 15



DEC MultiTitan

Fig. 17



Cray 1

Fig. 16

INTERNATIONAL SEARCH REPORT

International Application No

PC/GB 99/00707

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F15/78 G06F9/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No
Y	GB 2 216 307 A (ARDENT COMPUTER CORP) 4 October 1989 see page 7, line 7 - page 8, line 30 see page 13, line 7 - page 14, line 2 ---	1-15
Y	EP 0 646 877 A (FUJITSU LTD) 5 April 1995 see page 10, line 17 - line 29 ---	1-15
A	US 5 513 366 A (AGARWAL RAMESH C ET AL) 30 April 1996 see column 10, line 13 - line 61 see column 11, line 24 - column 12, line 13 see column 12, line 43 - line 60 -----	1,12



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claims or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, the combination being obvious to a person skilled in the art
- document member of the same patent family

Date of the actual completion of the international search

6 July 1999

Date of mailing of the international search report

13/07/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel: (+31-70) 340-2040, Tx: 31 651 epo nl.
Fax: (+31-70) 340-3016

Authorized officer

Daskalakis, T

INTERNATIONAL SEARCH REPORT

International application No.
PCT/GB 99/00707

Box I Observations where certain claims were found unsearchable (Continuation of item 1 of first sheet)

This International Search Report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the International Application that do not comply with the prescribed requirements to such an extent that no meaningful International Search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☐ As all required additional search fees were timely paid by the applicant, this International Search Report covers all searchable claims.
2. ☒ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this International Search Report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this International Search Report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest.
- ☐ No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. Claims: 1-7,9-15

Apparatus for processing data with register bank comprising a plurality of subsets of sequences of registers.

2. Claim : 8

Apparatus for processing data with register bank comprising at least one subset of registers and a transfer controller for controlling transfers of data between the memory and said register bank

INTERNATIONAL SEARCH REPORT

Information on patent family members

Inventor's International Application No

PCT/GB 99/00707

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
GB 2216307 A	04-10-1989	DE 3906327 A FR 2628237 A JP 2010467 A	14-09-1989 08-09-1989 16-01-1990
EP 0646877 A	05-04-1995	JP 7152733 A US 5669013 A	16-06-1995 16-09-1997
US 5513366 A	30-04-1996	NONE	